
Tensor Network Impurity Solvers

Simulating Quantum Materials

Martin Grundner



München 2025

Tensor Network Impurity Solvers

Simulating Quantum Materials

Martin Grundner

Dissertation
der Fakultät für Physik
der Ludwig-Maximilians-Universität
München

vorgelegt von
Martin Grundner
aus Graz

München, den 13.01.2025

Erstgutachter: Prof. Dr. Ulrich Schollwöck
Zweitgutachter: Prof. Dr. Christian Mendl
Tag der mündlichen Prüfung: 28.03.2025

Kurzfassung

Übergangsmetalloxide (TMOs) und andere *Quantenmaterialien* entfachen neuerlich großes Interesse aufgrund ihrer Vielzahl an funktionalen Eigenschaften. Diese entstehen typischerweise durch starke elektronische Korrelationseffekte bei niedrigen Temperaturen, was ihre theoretische Beschreibung äußerst komplex gestaltet. Die derzeit erfolgreichste Methode zur Simulation solcher Materialien ist die dynamische Molekularfeldtheorie (DMFT). Diese zeichnet sich durch eine exakte Beschreibung nichtlokaler kinetischer Terme aus, während die Coulomb-Wechselwirkung als lokal angenommen wird.

Wir präsentieren sowohl methodische als auch konzeptionelle Fortschritte, welche uns die Simulation von TMOs bei tiefen Temperaturen innerhalb der DMFT-Approximation mit Tensornetzwerk-basierten Algorithmen ermöglichen. Insbesondere führen wir eine Baum-Tensornetzwerkstruktur ein, den MT3N, welche speziell darauf ausgelegt ist, die komplexe Korrelationsstruktur von Multiorbital-Modellen bestmöglich darzustellen. Ein wesentlicher Vorteil unserer Tensornetzwerk-basierten Algorithmen ist die Möglichkeit Green'sche Funktionen sowohl auf der Matsubara Achse, als auch direkt auf der reellen Frequenzachse zu berechnen. Wir haben einen neuen Algorithmus zur analytischen Fortsetzung, MinKL, entwickelt, der es uns ermöglicht, Informationen beider Green'schen Funktionen zu kombinieren, wodurch das Verfahren signifikant stabilisiert und die Genauigkeit im Vergleich zu gängigen Algorithmen deutlich verbessert wird.

Zusätzlich stellen wir die Möglichkeit quantenmechanische Systeme entlang komplexer Zeitkonturen zu entwickeln vor. Dies führt dazu, dass hochenergetische Zustände mit zunehmenden Abstand zur reellen Zeitachse immer stärker unterdrückt werden. Dadurch kann das Verschränkungswachstum während der Zeitentwicklung signifikant reduziert werden, was zu erheblichen Laufzeit- und Genauigkeitsverbesserungen führt. Zudem diskutieren wir mehrere komplexe Zeitkonturen sowie verschiedene Methoden um Ergebnisse auf die reelle Frequenzachse analytisch fortzusetzen. Damit können wir das Fermi-Flüssigkeits-Verhalten von Multiorbital Systemen bis zu ≈ 0.002 eV auflösen.

Abschließend untersuchen wir das Übergangsmetalloxid LiV_2O_4 , welches schwere Quasiteilchen bei tiefen Temperaturen formiert, eine Eigenschaft, welche typischerweise nur in f-Orbital-Materialien auftritt. Unsere Entwicklungen ermöglichen es uns die Entstehung dieser Quasiteilchen akkurat zu simulieren und eine neue Theorie dahinter zu präsentieren.

Abstract

Transition metal oxides (TMOs) and other *quantum materials* recently attracted immense interest due to their plethora of functional properties. These properties are often the result of strong electronic correlation effects at low temperatures, rendering their theoretical description challenging. The current state-of-the-art method for the simulation of such materials is dynamical mean field theory (DMFT), which provides an exact description of the kinetic part of the system while approximating the Coulomb interaction as local.

In this thesis, we present advances in tensor network based impurity solvers, which we use to simulate intricate TMOs at low temperatures within the DMFT approximation. We discuss both methodological and conceptual developments that result in significant improvements in runtime and accuracy. We introduce a tree tensor network structure, the MT3N, specifically tailored to optimally represent the intricate correlation structure of multi-orbital impurity models. A significant advantage of tensor network based impurity solvers is their ability to compute Green's functions both on the Matsubara axis and directly on the real frequency axis. We developed a new analytic continuation algorithm, MinKL, that allows us to combine those Green's functions to significantly stabilize the procedure and improve its accuracy compared to prevalent algorithms.

Additionally, we introduce a novel concept of time evolution by evolving systems along complex time contours. By shifting time evolution away from the real time axis, we significantly curtail entanglement growth, enabling substantial improvements in accuracy and efficiency. We present several complex time contours along with multiple post-processing methods that analytically continue our results back to the real frequency axis. These advancements enable us to resolve the Fermi liquid behavior of a multi-orbital system down to ≈ 0.002 eV.

Finally, we use these developments in the study of the transition metal oxide LiV_2O_4 . This material has captivated researchers due to its heavy quasiparticle mass at low temperatures, a rare occurrence outside f-orbital materials. Our algorithmic advancements allow us to propose a new theory describing this compound's emerging heavy fermion regime.

List of Publications

This dissertation is partially based on the author's publications in peer-reviewed journals [1–6] and preprint servers [7, 8]. This list is sorted in inverse chronological order.

LiV₂O₄: Hund-Assisted Orbital-Selective Mottness

M. Grundner, F.B. Kugler, O. Parcollet, U. Schollwöck, A. Georges and A. Hampel
<https://doi.org/10.48550/arXiv.2409.17268>

Complex time evolution in tensor networks and time-dependent Green's functions

M. Grundner, P. Westhoff, F. B. Kugler, O. Parcollet and U. Schollwöck
<https://link.aps.org/doi/10.1103/PhysRevB.109.155124>

Cooper-paired bipolaronic superconductors

M. Grundner, T. Blatz, J. Sous, U. Schollwöck and S. Paeckel
<https://doi.org/10.48550/arXiv.2308.13427>

Stable bipolarons in open quantum systems

M. Moroder, **M. Grundner**, F. Damanet, U. Schollwöck, S. Mardazad, S. Flannigan, T. Köhler, and S. Paeckel
<https://doi.org/10.1103/PhysRevB.107.214310>

Formation of CuO₂ sublattices by suppression of interlattice correlations in tetragonal CuO

M. Bramberger, B. Bacq-Labreuil, **M. Grundner**, S. Biermann, U. Schollwöck, S. Paeckel, and B. Lenz
<https://www.scipost.org/SciPostPhys.14.1.010>

Quantum dynamics simulation of intramolecular singlet fission in covalently linked tetracene dimer

S. Mardazad, Y. Xu, X. Yang, **M. Grundner**, U. Schollwöck, H. Ma, and S. Paeckel
<https://doi.org/10.1063/5.0068292>

BaOsO₃: A Hund's metal in the presence of strong spin-orbit coupling

M. Bramberger, J. Mravlje, **M. Grundner**, U. Schollwöck, and M. Zingl
<https://link.aps.org/doi/10.1103/PhysRevB.103.165133>

Sr₂MoO₄ and Sr₂RuO₄: Disentangling the roles of Hund's and van Hove physics

J. Karp, M. Bramberger, **M. Grundner**, U. Schollwöck, A. J. Millis, and M. Zingl
<https://link.aps.org/doi/10.1103/PhysRevLett.125.166401>

Contents

Zusammenfassung	iii
Abstract	iv
List of Publications	vi
1 Introduction	1
2 Tensor Networks	5
2.1 Matrix Product States	6
2.1.1 Decomposition	7
2.1.2 Tensor Diagrams	13
2.2 Density Matrix Renormalization Group	14
2.3 Time-Dependent Variational Principle (TDVP)	18
2.3.1 Global Subspace Expansion (GSE)	22
2.3.2 Local Subspace Expansion (LSE)	24
3 Dynamical Mean Field Theory	27
3.1 Impurity Models	30
3.1.1 Real Material Simulations	32
3.2 Impurity Solvers	33
4 Imaginary Time Impurity Solver	37
4.1 Discretization	37
4.2 Ground state search	38
4.3 Time Evolution	40
4.3.1 Linear Prediction	46
4.3.2 How to Perform Efficient Simulations	49
4.4 Analytic Continuation	51
4.4.1 MinKL - High Precision Analytic Continuation	55
5 Real Time Impurity Solver	63
5.1 Discretization	63
5.1.1 Improved Estimators	73

5.2	Optimal Basis	75
5.3	Ground State Search	85
5.4	Time evolution	88
6	Multi-Orbital Impurity Solvers	91
6.1	Three-legged Tensor Network States	92
6.2	Comparison with state-of-the-art Tensor networks	96
7	Complex Time Impurity Solver	101
7.1	Constant Imaginary Time: Parallel Contour	104
7.2	Time Evolution at Fixed Angles: Tilted Contour	110
7.3	Real Frequency Results Without Analytic Continuation	113
7.4	Multi-Orbital Problems	115
7.5	Summary	119
8	LiV₂O₄ - A Heavy Fermion TMO	121
8.1	Material	122
8.2	DFT+DMFT	125
8.3	Purification	127
	8.3.1 How to stabilize finite temperature simulations	129
	8.3.2 Odd/Even effect	132
8.4	Results	134
8.5	Summary	140
9	Conclusion and Outlook	141
A	Post-processing for Complex Time Contours	145
A.1	Extrapolation Method	146
A.2	Stabilizing the Kink Method	147
B	LiV₂O₄: Comparison of Multi-orbital models	149
	Acknowledgements	172

Chapter 1

Introduction

The development of the band theory of solids was fundamental to the digital revolution of the 20th century. It provides a framework for understanding how valence electrons are organized within a solid and how they contribute to its properties. When atomic orbitals of solids are strongly overlapping, their electronic structure is typically well described by an approximation of nearly free electrons where nuclei are assumed to be stationary, generating a periodic background potential. Instead of discrete energy levels, as in isolated atoms, electronic energy levels in solids form a smooth band in momentum space [9]. This allowed for a simple classification of materials into metals and insulators, i.e., a metal is a solid with a partially filled band, whereas insulators are materials in which the Fermi level, the energy up to which bands are filled, lies between two bands. Insulators can be further subdivided into insulators and semiconductors, where the latter are typically characterized by a smaller band gap. Bands in semiconductors can be shifted towards the Fermi level upon appropriate doping, thus significantly increasing the number of partially filled states. Such specifically tuned semiconductors are at the heart of modern field effect transistors [10]; and by that, they are in the pocket of nearly every person on earth.

Modern material research is now, again, at the forefront of scientific progress driven by an ever-growing demand for higher efficiency [11]. This led to the emergence of synthetic *quantum materials* with tailored electronic properties, determined by quantum effects near the Fermi edge [12, 13]. Promising candidates for such engineered materials are transition metal oxides (TMOs), where strong electronic correlation effects give rise to a plethora of functional properties, including superconductivity, Mott transitions, multiferroicity, and strange metallicity [14–16]. Advancements in material synthesis techniques, especially those involving thin-film deposition [17] and molecular-beam epitaxy [18], have enabled the creation of layered artificial materials with precise control over their electronic properties [14, 15]. However, the theoretical description of such materials is complicated. Typically, TMOs have partially filled bands that stem from d-orbitals, which are spatially strongly confined. This leads to a competition between kinetic energy terms of single electrons and interaction effects between them. Thus, a simple description of effective potentials in what is commonly referred to as *strongly correlated materials* becomes insufficient [19].

This includes the simple description of nearly free electrons as well as more advanced approaches like density functional theory (DFT) [20–22] that considers the electron–electron interactions via a self-consistently obtained effective potential [23].

An exact quantitative description of *strongly correlated materials* is numerically unfeasible. However, in the past years, immense progress has been made in combining single-particle approaches, like DFT, with more elaborate methods that directly account for interaction effects [24–26]. Most notably, the combination of DFT with dynamical mean field theory (DMFT) [27, 28], which is typically referred to as (DFT + DMFT), established itself as the state-of-the-art for electronic structure simulations of *strongly correlated materials* [1, 2, 29–33]. DMFT transforms the interacting lattice model into an impurity problem by considering electron–electron interactions solely within the impurity cluster. The dynamic exchange between this cluster and the remaining lattice is represented through a self-consistently optimized bath of free electrons [27, 29, 34]. Although this constitutes a significant simplification of the problem at hand, it is nonetheless computationally demanding as it needs to be solved for all interacting single-particle Green’s functions of the impurity cluster. As such, the actual solution of the impurity problem has evolved into its own field of research, and large efforts have been devoted to developing efficient impurity solvers.

In this thesis, we developed multiple tensor network based impurity solvers that work both on the imaginary frequency axis [1, 2, 33–35] as well as directly on the real frequency axis [6, 36–39]. These solvers have significant advantages over commonly used solvers. They are neither restricted by the system size, unlike exact diagonalization (ED) [40, 41], nor the specific type of interactions present in the impurity cluster, unlike the numerical renormalization group (NRG) [42–44] or the presence of a sign problem, unlike continuous time quantum Monte Carlo (CTQMC) [45]. We showcase the potential of our imaginary time solver on the basis of the heavy fermion TMO LiV_2O_4 [46], which experiences dramatic electronic correlations at very low temperatures, inaccessible to CTQMC. We precede our presentation of these results with a complete discussion of our impurity solvers and algorithmic advances therein.

Additionally, we introduce a novel approach in the realm of real frequency solvers: complex time evolution [6]. Shifting time evolution contours away from the real time axis significantly curtails entanglement growth, enabling substantial improvements in accuracy and efficiency. This is due to the additional time evolution towards the imaginary time direction, which effectively projects the state into a low-energy subspace by progressively dampening high-energy contributions. We present several analytic continuation procedures that continue our results back to the real frequency axis. Among them are simple analytic expressions, rendering this approach stable, fast, and highly precise.

This thesis is structured as follows: Chapter 2 provides an overview of tensor network methods, with a particular focus on matrix product states (MPS) and tree tensor network states (TTNS). In chapter 3, we explore dynamical mean field theory and briefly analyze typical impurity Hamiltonians used in realistic material simulations. Chapter 4 details our development of an imaginary time impurity solver, highlighting our implemented enhancements. Similarly, chapter 5 discusses our advancements in the domain of real frequency impurity solvers. In chapter 6, we introduce innovative tensor network structures tailored for multiorbital impurity models. Our novel complex time impurity solver is presented in chapter 7, where it is compared against our traditional real time solver and numerical renormalization group (NRG) methods. Chapter 8 delves into our DFT+DMFT study of the heavy fermion TMO compound LiV_2O_4 . The thesis concludes with chapter 9, summarizing our principal discoveries and contributions to the field.

Chapter 2

Tensor Networks

Tensor networks (TN) represent a broad category of sophisticated yet intuitively appealing approaches for decomposing quantum many-body systems. Central to these methods are rank-revealing decompositions, especially the singular value decomposition (SVD). These methods have recently garnered increased interest and application across mathematics and various physics disciplines. In the mathematical community, these approaches are, in their one-dimensional variant, referred to as tensor trains (TT) [47]. The advent of new decomposition strategies, like the tensor cross interpolation (TCI) algorithm [48, 49], has broadened the application of these concepts to fields such as numerical integration [50, 51] and, consequently, Monte Carlo methods [52]. It also sparked a dynamic new area of research known as quantized tensor trains (QTT) [49, 53]. QTTs are pivotal in achieving highly efficient representations of large-scale, structured data, including correlation functions [53, 54] and Brillouin zone integrals [55]. Beyond these applications, tensor network methods have been applied in modeling large biological systems [56], simulating turbulent flows [57, 58], enhancing reinforcement learning algorithms [59] and compressing large language models [60].

This chapter provides a foundational overview of tensor networks. Given their critical role in the simulation of dynamic quantities like impurity Green's functions, particular emphasis will be given to time evolution methods for tensor networks. We will explore various methodological advancements and present a revised "blueprint" for time evolution within the framework of tensor networks. We aim to highlight the historical progression of tensor network methods, anchoring our discussion in the foundational concepts of Matrix Product States (MPS). As a one-dimensional tensor network state, MPS remains the most commonly used tensor network state due to its many favorable numerical properties.

2.1 Matrix Product States

Low-dimensional *strongly correlated* quantum systems; chains of interacting electrons or electron-phonon systems, to name but a few, are host to a myriad of interesting correlation effects. However, exact analytical solutions for these complex many-body systems are scarce. The Bethe Ansatz [61, 62], a sophisticated method for solving integrable systems primarily in one dimension, has been applied to certain cases. However, more complex models or realistic materials are generally dependent on numerical techniques.

In 1975, Kenneth Wilson revolutionized the field of condensed matter physics by developing the numerical renormalization group (NRG) method. This groundbreaking work gave a precise description of the Kondo model's ground state, elucidating how conduction electrons interact with magnetic impurities [42]. Wilson's contribution extended beyond the introduction of a mere algorithm; he highlighted the significance and effectiveness of the renormalization group (RG) approach. This concept has since been embraced and refined across various disciplines in physics, culminating in Wilson's recognition with the Nobel Prize. The core principle of RG methods lies in isolating the pertinent degrees of freedom. Wilson achieved this by applying a logarithmic discretization to the conduction band, crafting an effective Hamiltonian that adeptly captured the low-energy dynamics by diminishing the influence of high-energy states. Despite its innovative approach, the application of NRG has been primarily confined to impurity models since its inception, with limited success in broader contexts [63].

Steven White's development of the density matrix renormalization group (DMRG) in 1992 marked a significant breakthrough in the study of low-dimensional quantum systems [63, 64]. Quickly emerging as the leading numerical method for analyzing strongly correlated one-dimensional systems, DMRG excelled in both ground state simulations and the investigation of dynamical properties [65–67]. Unlike NRG, which focuses on energy-based state reduction, White's method innovatively prioritizes states based on their contribution to the density matrix, enhancing efficiency and accuracy. We aim to delve into White's rationale behind the asserted optimality of this truncation approach. Additional justifications are provided in [67]. We will closely follow [67, 68] in our derivation.

Consider that $|\psi\rangle$ is the state of a bipartite system $AB = A \otimes B$ given as

$$|\psi\rangle = \sum_{a=1}^{m_A} \sum_{b=1}^{m_B} c_{ab} |a\rangle |b\rangle, \quad (2.1)$$

where $c_{ab} \in \mathbb{C}$ are wave function coefficients, $|a\rangle, |b\rangle$ are basis states in A and B respectively and m_a, m_b are the number of respective basis states. Let us now construct an approximation on above wave function $|\tilde{\psi}\rangle$ by keeping fewer basis states in A ,

$$|\tilde{\psi}\rangle = \sum_{a'=1}^{m'_A} \sum_{b=1}^{m_B} \tilde{c}_{a'b} |a'\rangle |b\rangle, \quad m'_A \leq m_A, \quad (2.2)$$

with $|a'\rangle$ denoting the new basis states. The optimal coefficients $\tilde{c}_{a'b} |a'\rangle$ are determined by minimizing the quadratic norm $\| |\psi\rangle - |\tilde{\psi}\rangle \|_2^2$. This expression is minimized by selecting the m'_A eigenvectors $|a'\rangle$ corresponding to the largest eigenvalues $w_{a'}$ of the density matrix $\rho_{aa'} = \sum_b c_{ab} c_{a'b}$ [63]. We can thus measure the error of this truncation procedure as

$$\| |\psi\rangle - |\tilde{\psi}\rangle \|_2^2 = 1 - \sum_{a'}^{m'_A} w_{a'} = \epsilon_t. \quad (2.3)$$

In tensor network algorithms, these eigenvectors are typically obtained using singular value decompositions [63].

With a truncation procedure in place, we now focus on the concept of matrix product states (MPS). Due to their historical development, DMRG and MPS are often used interchangeably. However, in the wake of advancements in complex tensor network structures [69–71], a distinction is commonly made: the decomposed state is known as a matrix product state, while the algorithmic approach to finding the ground state is referred to as DMRG.

2.1.1 Decomposition

Although matrix product states inherently resemble a one-dimensional chain, their application extends beyond simulating one-dimensional systems [72]. For the sake of clarity, however, let us consider a one-dimensional lattice system comprising L sites. An arbitrary pure state on this lattice can be described as follows:

$$|\psi\rangle = \sum_{\{\sigma_i\}} c_{\sigma_1 \dots \sigma_L} |\sigma_1, \dots, \sigma_L\rangle, \quad (2.4)$$

where $\{\sigma_i\} = \sigma_1, \dots, \sigma_L$ denotes the set of d_i basis states of the local Hilbert space \mathcal{H}_i , and $c_{\sigma_1 \dots \sigma_L}$ is the L -dimensional coefficient tensor. In general, the coefficient tensor's complexity increases exponentially with the size of the system. In the exceptional uncorrelated cases where the quantum state is representable as a product state, this tensor simplifies, allowing decomposition into a product of individual components $c_{\sigma_1 \dots \sigma_L} = \prod_i c_{\sigma_i}$. Nonetheless, managing this tensor's complexity remains a formidable challenge for systems with strong electron correlations.

Although a trivial decomposition is not possible in the presence of entanglement, we can construct a decomposition scheme that still preserves the correlation of the state while leaving us with a local, decomposed representation of the coefficient tensor

$$c_{\sigma_1 \dots \sigma_L} = \sum_{m_1, \dots, m_{L-1}} M_{1, m_1}^{\sigma_1} M_{m_1, m_2}^{\sigma_2} \dots M_{m_{L-1}, 1}^{\sigma_L}, \quad (2.5)$$

with virtual bond indices m_i . Contracting identically named bond indices results in a matrix multiplication of the respective site tensor M . Consequently, the entire coefficient

tensor can be reconstructed via successive matrix multiplications, thus marking the origin of the term 'matrix product state'. This decomposition can be obtained by subsequent singular value decompositions of the coefficient tensor. Let us showcase this idea on the example of what is referred to as a *left-canonical* MPS. We follow a derivation presented in [73].

Let us reshape the coefficient tensor as $c_{\sigma_1 \dots \sigma_L} = \Psi_{\sigma_1, (\sigma_2 \dots \sigma_L)}$ where $\Psi_{\sigma_1, (\sigma_2 \dots \sigma_L)}$ is a $(d \times d^{L-1})$ dimensional matrix. A singular value decomposition of Ψ gives

$$\Psi_{\sigma_1, (\sigma_2 \dots \sigma_L)} = \sum_{a_1}^{r_1} U_{\sigma_1, a_1} S_{a_1, a_1} V_{a_1, (\sigma_2 \dots \sigma_L)}^\dagger, \quad (2.6)$$

with rank, also commonly called bond dimension, $r_1 \leq d$. We can multiply S and V^\dagger and reshape the resulting matrix into a vector $c_{a_1 \dots \sigma_L}$. Decomposing the U matrix into a collection of d row vectors A^{σ_1} with entries $A_{a_1}^{\sigma_1} = U_{\sigma_1, a_1}$ gives us

$$c_{\sigma_1 \dots \sigma_L} = \Psi_{\sigma_1, (\sigma_2 \dots \sigma_L)} = \sum_{a_1}^{r_1} A_{a_1}^{\sigma_1} c_{a_1 \sigma_2 \dots \sigma_L}. \quad (2.7)$$

Reshaping $c_{a_1 \sigma_2 \dots \sigma_L}$ and subsequently performing the SVD as presented in equation (2.6) gives

$$c_{\sigma_1 \dots \sigma_L} = \Psi_{\sigma_1, (\sigma_2 \dots \sigma_L)} = \sum_{a_1}^{r_1} A_{a_1}^{\sigma_1} c_{a_1 \dots \sigma_L} = \sum_{a_1}^{r_1} \sum_{a_2}^{r_2} A_{a_1}^{\sigma_1} A_{a_1, a_2}^{\sigma_2} c_{a_2 \sigma_3 \dots \sigma_L}, \quad (2.8)$$

with rank $r_2 \leq r_1 d \leq d^2$. Iteratively performing this procedure leaves us with this decomposition

$$c_{\sigma_1 \dots \sigma_L} = \sum_{a_1, \dots, a_{L-1}} A_{a_0, a_1}^{\sigma_1} A_{a_1, a_2}^{\sigma_2} \dots A_{a_{L-2}, a_{L-1}}^{\sigma_{L-1}} A_{a_{L-1}, a_L}^{\sigma_L}, \quad (2.9)$$

where we added dummy indices $a_0 = 1$ and $a_L = 1$ to consistently denote A as matrices. The rank of a_0 or a_L can also be larger than one, for example, if one wants to express sums of states in different quantum number sectors. This finally leaves us with a definition of a *left-canonical* matrix product state:

$$|\psi\rangle = \sum_{\{\sigma_i\}} A^{\sigma_1} \dots A^{\sigma_L} |\sigma_1, \dots, \sigma_L\rangle. \quad (2.10)$$

The name "left-canonical" refers not solely to the iterative construction principle, see figure 2.1, but also to special properties gained due to the iterative use of singular value decompositions. As $U^\dagger U = \mathbb{1}$ holds for U matrices that are the outcome of a SVD, the following relation holds for the tensors obtained by the procedure above [73]:

$$\sum_{\sigma_i} A^{\sigma_i \dagger} A^{\sigma_i} = \mathbb{1}. \quad (2.11)$$




Figure 2.1: Decomposition of a quantum state into a matrix product state by the iterative application of singular value decompositions. Grey bonds mark the outermost virtual legs with bond dimension 1. The coefficient tensor can be restored by contracting over all virtual legs (horizontal bonds).

We refer to A matrices fulfilling this property as left-normalized. As the choice to start the decomposition procedure from the left is arbitrary, we can also perform a similar decomposition strategy from the right and obtain the *right-canonical* form

$$|\psi\rangle = \sum_{\{\sigma_i\}} B^{\sigma_1} \dots B^{\sigma_L} |\sigma_1, \dots, \sigma_L\rangle, \quad (2.12)$$

where, again, exploiting properties of the SVD, $VV^\dagger = \mathbb{1}$ we find that B matrices obey

$$\sum_{\sigma_i} B^{\sigma_i} B^{\sigma_i\dagger} = \mathbb{1}. \quad (2.13)$$

As alluded to in our initial motivation for tensor networks, much of its success is rooted in a truncation procedure based on Schmidt values. However, neither a bipartition of the *left canonical* nor the *right canonical* form would result in a proper Schmidt decomposition [73], thus significantly limiting the efficiency of local optimization procedures. Let us consider this at the example of a bipartition of a *left-canonical* state $|\psi\rangle$, as presented in [73], in two subsystems A and B:




Figure 2.2: Mixed canonical MPS. The first four sites are left normalized, S_4 denotes the singular value matrix, and the four rightmost sites are right normalized.

$$|a_i\rangle_A = \sum_{\sigma_1, \dots, \sigma_i} (A^{\sigma_1} \dots A^{\sigma_i})_{1, a_i} |\sigma_1, \dots, \sigma_i\rangle, \quad (2.14)$$

$$|a_i\rangle_B = \sum_{\sigma_{i+1}, \dots, \sigma_L} (A^{\sigma_{i+1}} \dots A^{\sigma_L})_{a_i, 1} |\sigma_{i+1}, \dots, \sigma_L\rangle, \quad (2.15)$$

which allows us to write the total state as

$$|\psi\rangle = \sum_{a_i} |a_i\rangle_A |a_i\rangle_B. \quad (2.16)$$

However, this does not represent a Schmidt decomposition as $|a_i\rangle_B$ does not form an orthonormal set in general. For $|a_i\rangle_A$ this property immediately follows from their construction, we find

$${}_A \langle a'_i | a_i \rangle_A = \sum_{\sigma_1, \dots, \sigma_i} (A^{\sigma_1} \dots A^{\sigma_i})_{1, a'_i}^* (A^{\sigma_1} \dots A^{\sigma_i})_{1, a_i}, \quad (2.17)$$

$$= \sum_{\sigma_1, \dots, \sigma_i} (A^{\sigma_1} \dots A^{\sigma_i})_{a'_i, 1}^\dagger (A^{\sigma_1} \dots A^{\sigma_i})_{1, a_i}, \quad (2.18)$$

$$= \sum_{\sigma_1, \dots, \sigma_i} (A^{\sigma_i^\dagger} \dots \underbrace{A^{\sigma_1^\dagger} A^{\sigma_1}}_{\delta_{a'_1, a_1}} \dots A^{\sigma_i})_{a'_i, a_i}, \quad (2.19)$$

$$= \delta_{a'_i, a_i}. \quad (2.20)$$

In the last step, we iteratively used the left normalization defined in equation (2.11). Calculating the overlap in B yields

$${}_B \langle a'_i | a_i \rangle_B = \sum_{\sigma_{i+1}, \dots, \sigma_L} (A^{\sigma_{i+1}} \dots A^{\sigma_L})_{a'_i, 1}^* (A^{\sigma_{i+1}} \dots A^{\sigma_L})_{a_i, 1}, \quad (2.21)$$

$$= \sum_{\sigma_{i+1}, \dots, \sigma_L} (A^{\sigma_{i+1}} \dots A^{\sigma_L} A^{\sigma_L^\dagger} \dots A^{\sigma_{i+1}^\dagger})_{a'_i, a_i}, \quad (2.22)$$

which, in general, cannot be simplified further.




Figure 2.3: Mixed canonical MPS with respect to an active site. The first three sites are left normalized, the center site denotes the active site, and the four rightmost sites are right normalized.

A similar problem arises for right-canonical states with the difference that states in A will generally not form an orthonormal basis [73]. Hence, most local tensor network algorithms use a further canonical form, the *mixed-canonical* representation. Assume we perform a decomposition from the left up to site i as

$$c_{\sigma_1 \dots \sigma_L} = \sum_{a_i} (A^{\sigma_1} \dots A^{\sigma_i})_{a_i} S_{a_i, a_i} V_{a_i, (\sigma_{i+1} \dots \sigma_L)}^\dagger, \quad (2.23)$$

and $V_{a_i, (\sigma_{i+1} \dots \sigma_L)}^\dagger$ can be subsequently decomposed from the right to [73]

$$|\psi\rangle = \sum_{\sigma_1, \dots, \sigma_L} A^{\sigma_1} \dots A^{\sigma_i} S B^{\sigma_{i+1}} \dots B^{\sigma_L} |\sigma_1, \dots, \sigma_L\rangle, \quad (2.24)$$

where S is a diagonal matrix containing the singular values of the bond between i and $i+1$, see figure 2.2.

It becomes immediately apparent that by construction, this leads to a Schmidt decomposition upon bipartition of the state into subsystems A and B . Let us, again, consider such a bipartition, now given as

$$|a_i\rangle_A = \sum_{\sigma_1, \dots, \sigma_i} (A^{\sigma_1} \dots A^{\sigma_i})_{1, a_i} |\sigma_1, \dots, \sigma_i\rangle \quad (2.25)$$

$$|a_i\rangle_B = \sum_{\sigma_{i+1}, \dots, \sigma_L} (B^{\sigma_{i+1}} \dots B^{\sigma_L})_{a_i, 1} |\sigma_{i+1}, \dots, \sigma_L\rangle \quad (2.26)$$

which leads to the Schmidt decomposition

$$|\psi\rangle = \sum_{a_i} S_{a_i, a_i} |a_i\rangle_A |a_i\rangle_B. \quad (2.27)$$

The mixed canonical representation is the representation of choice for local tensor network algorithms due to its efficient representation of inactive sites (sites that are not currently updated) and its ability to generate a Schmidt decomposition upon bipartition.




Figure 2.4: Maximum bond dimension profile of a MPS. The highest bond dimension in an MPS of length 8 is $d_{\max} = d^{\frac{L}{2}} = d^4$.

This comes as a consequence of the left and right normality of the subsystems. Note that this representation is typically used in respect to what is called an active site

$$|\psi\rangle = \sum_{\sigma_1, \dots, \sigma_L} A^{\sigma_1} \dots A^{\sigma_{i-1}} M^{\sigma_i} B^{\sigma_{i+1}} \dots B^{\sigma_L} |\sigma_1, \dots, \sigma_L\rangle, \quad (2.28)$$

where the singular value tensor from equation (2.24) has been multiplied to the left, see figure 2.3.

Equipped with a construction principle for matrix product state representations, let us shift the focus from how to why; why are matrix product state decompositions efficient? Let us reconsider the bond dimension profile of matrix product states. As we have seen in equation (2.6), the bond dimension on the right virtual leg $m_{i+1} = m_i d_i$ is upper bounded by the left virtual leg, and the physical dimension of the site. Assuming the maximum bond dimension for a complete MPS, this leaves us with

$$(1, d), (d, d^2), \dots, (d^{L/2-1}, d^{L/2}), (d^{L/2}, d^{L/2-1}), \dots, (d^2, d), (d, 1), \quad (2.29)$$

where the numbers in brackets refer to the virtual leg on the left and right, respectively. For simplicity, we assumed an even number of sites, the bond dimension of the MPS is upper bounded by $d^{L/2}$ as the upper bound on the bond dimension of a site tensor must be fulfilled from the left as well as the right of the system, see figure 2.4 for a visualization.

In a simulation of $\mathcal{O}(1000)$ sites, this exact decomposition is just as impractical as exact diagonalization. Fortunately, in practical applications, the theoretical upper bounds on bond dimension remain often unattained, thus permitting the construction of efficient yet numerically exact tensor network algorithms. An explanation for this can be found in the *area law* of entanglement, which dictates that the growth of entanglement between a subsystem and the finite remainder of the system is at most proportional to their boundary [74, 75].

Specifically, for one-dimensional systems, this implies that the entanglement across any bipartition would plateau, as the boundary of a one-dimensional system is constant. While the area law does not hold in general, it has been rigorously validated for ground states of local gapped Hamiltonians in one dimension [74–76]. This is why matrix product states reign nearly uncontested as the numerical method of choice for one-dimensional quantum systems.





Figure 2.5: Tensor leg direction convention for matrix product states.

2.1.2 Tensor Diagrams

Before we move on to the formulation of tensor network algorithms, let us properly introduce the commonly used short-hand graphical notation for tensors. Consider a tensor T as a map from input vector spaces A and B to output vector spaces C and D and let us further construct a basis in these vector spaces as $\{|a\rangle_A\}$, $\{|b\rangle_B\}$, $\{|c\rangle_C\}$ and $\{|d\rangle_D\}$. We can define this map with tensor coefficients T_{ab}^{cd} as:

$$T : A \otimes B \rightarrow C \otimes D, \quad (2.30)$$

$$T : |a\rangle_A \otimes |b\rangle_B \mapsto |c\rangle_C \otimes |d\rangle_D = \sum_{ab} T_{ab}^{cd} |a\rangle_A \otimes |b\rangle_B, \quad (2.31)$$

Figure 2.6: Graphic representation of the tensor defined in equation (2.30). States from input/output spaces are indicated by arrows pointing to/from the tensor T .

We represent this rank four tensor in graphic notation as a circle with four legs, each corresponding to one of the four vector spaces with arrows indicating the direction. So far, we deliberately disregarded the direction of our tensor legs, as, for matrix product states, there is some freedom in how one chooses to define them. In our tensor network toolkit, SYTEN [77, 78] a tensor $M_{m_{i-1}, m_i}^{\sigma_i}$ is to be understood as $M_{\sigma_i m_i}^{m_{i-1}}$ where lower indices denote incoming legs and upper indices denote outgoing legs.

For graphical representations of large tensor networks, we will continue to disregard leg directions for simplicity, but we will present them for tensor contractions in our discussion of algorithms. A graphical representation of an MPS with this convention can be found in figure 2.5.




Figure 2.7: Quantum number flow in matrix product states. The right dummy leg (m_L) represents the vacuum, while the left dummy leg (m_0) denotes the quantum number sector of the state.

The orientation of tensor legs acquires additional significance in the context of symmetries, interpreted as a "flow of quantum numbers" within the matrix product state framework [78]. By imposing the condition that the aggregate quantum number sectors of the incoming legs must match the quantum number sector of the outgoing leg, the right dummy leg (m_L) is naturally designated as the vacuum leg. Conversely, the left dummy leg (m_0) represents the quantum number sector of the state, as illustrated in figure 2.7.

2.2 Density Matrix Renormalization Group

This section aims to provide a concise overview of the density matrix renormalization group (DMRG); a variational optimization algorithm for determining the ground states of quantum many-body systems. More extensive reviews on DMRG can be found in [67, 73], for reviews of DMRG adjacent algorithms like the infinite size density matrix renormalization group (iDMRG), see [79].

DMRG seeks to find a state $|\psi\rangle$ that minimizes the energy

$$\min_{|\psi\rangle} E, \quad \text{with} \quad E = \frac{\langle \psi | \hat{H} | \psi \rangle}{\langle \psi | \psi \rangle} \quad (2.32)$$

for a Hamiltonian \hat{H} where \hat{H} is a hermitian operator. If $|\psi\rangle$ is normalized, id est $\langle \psi | \psi \rangle = 1$, this is equivalent to minimizing $\langle \psi | \hat{H} | \psi \rangle$. Hence, we can rewrite the above equation using a Lagrange multiplier λ as

$$\min_{|\psi\rangle} \mathcal{E} = \min_{|\psi\rangle} \left(\langle \psi | \hat{H} | \psi \rangle - \lambda (\langle \psi | \psi \rangle - 1) \right). \quad (2.33)$$




Figure 2.8: Effective Hamiltonian at site $i = 3$. Black boxes indicated contractions of tensors required to build left (L_i) and right (R_i) environment tensors.

This approach outlines a valid minimization process, yet the full potential of the MPS-MPO framework is truly harnessed upon transitioning to a local optimization strategy. Given that the complexity of updates in this minimization scales exponentially with the number of sites involved, minimizing the number of active sites is generally advantageous. Typically, optimization procedures are defined in sweeping patterns, iteratively updating one or two sites at a time. Correspondingly, these algorithms are dubbed single-site DMRG or two-site DMRG respectively [80]. Formulations defining the optimization procedure over bond tensors, effectively describing zero-site introduced schemes, have been defined in [81].

For the sake of simplicity, let us introduce the single-site variant of DMRG. Consider a MPS in mixed canonical representation with respect to a single site M_i . We can find the extremal point of \mathcal{E} as

$$\frac{\partial \mathcal{E}}{\partial M_i^\dagger} = H_i^{\text{eff}} M_i - \lambda M_i \stackrel{!}{=} 0, \quad (2.34)$$

where H_i^{eff} is the effective Hamiltonian at site i , a rank 6 tensor constructed via the contraction of left and right environment tensors L_{i-1} and R_{i+1} with the MPO site tensor W_i , see figure 2.8.




Figure 2.9: Graphical representation of the single-site DMRG eigenvalue equation.

We can now identify equation (2.34) as a simple eigenvalue equation

$$H_i^{\text{eff}} M_i = \lambda M_i. \quad (2.35)$$

Determining the lowest eigenvalue and its corresponding eigenvector from the equation above results in a new, locally optimal MPS tensor, denoted as \tilde{M}_i . To closely approximate the true ground state, one must iteratively sweep through the system, back and forth, by first shifting the center site of the mixed canonical representation in a sweeping direction and subsequently performing the update as described in equation (2.35) until convergence is reached. The chosen bond dimension during the simulation plays a critical role in determining the fidelity of the MPS approximation to the true ground state. Without truncation, equation (2.35) would already constitute the global groundstate. So far, as the bond dimension is not altered, or at least not increased, during our update procedure the primary means of enhancing the MPS approximation quality would be starting with an initial state characterized by a higher bond dimension in the DMRG algorithm. Considering the impracticality of this approach, it is worth exploring strategies to dynamically increase the bond dimension within MPS.

An equivalent formulation to equation (2.35) can be derived for two sites contracted to a new site tensor in mixed canonical representation $M_{i,i+1}$. The application of the corresponding effective Hamiltonian and the subsequent singular value decomposition, used to separate $\tilde{M}_{i,i+1}$ into \tilde{A}_i and \tilde{M}_{i+1} , leads to a potential increase in bond dimension between sites i and $i+1$ by the local bond dimension $\tilde{m} = dm$ [73]. However, applications of the effective Hamiltonian to the site tensor scale as $\mathcal{O}(2m^3d^2w + m^2d^3w^2)$ which, compared to the scaling of the single sites scheme, leads to an additional factor d in both terms [78]. Hence, this constitutes a more expensive optimization procedure, especially for bosonic systems or other systems with large local Hilbert spaces [3, 5, 7, 82].




Figure 2.10: Expansion tensor used in DMRG3S and the local subspace expansion method (LSE) for a left-to-right sweep, see section 2.3.2. The expansion is performed after the optimization/evolution and before the shift of the orthogonality center/back-evolution for DMRG3S or LSE, respectively.

A dynamic growth of bond dimension, and a change in basis states, in favorably scaling single-site algorithms, can be achieved by the use of subspace expansion techniques [80, 83]. We want to introduce the concept of subspace expansion based on the work of Hubig et al. [83]. While there have been earlier attempts at combining expansion terms with single-site DMRG [80], they were hindered by poorly scaling perturbation terms that limited their applicability [78]. Subspace expansion terms and their combination with single-site DMRG have been discussed in the mathematics community as alternating minimum energy algorithms (AMEn) [84, 85], whereas the method introduced by Hubig et al. became known as strictly single-site DMRG (DMRG3S).

The update scheme of DMRG3S proceeds according to the original single-site algorithm with an additional expansion after the update step, before shifting the orthogonality center to the next site [83]. An expansion over the bond m_i that leaves the overall state invariant can be defined as

$$M_{i;\sigma_i} \rightarrow \tilde{M}_{i;\sigma_i} = \begin{bmatrix} M_{i;\sigma_i} & E_{i;\sigma_i} \end{bmatrix} \quad (2.36)$$

$$M_{i+1;\sigma_{i+1}} \rightarrow \tilde{M}_{i+1;\sigma_{i+1}} = \begin{bmatrix} M_{i+1;\sigma_{i+1}} \\ 0 \end{bmatrix}, \quad (2.37)$$

where $M_{i;\sigma_i}$ and $E_{i;\sigma_i}$ are matrices in the MPS bond dimensions. It is straightforward to see that the product of these matrices would leave the state invariant. Hubig et al. [83] proposed an expansion tensor of the form

$$E_i = \alpha L_{i-1} M_i W_i, \quad (2.38)$$

where all possible indices are contracted, and the right MPO and MPS legs are contracted with a split tensor S to yield a rank 3 expansion tensor, as shown in figure 2.10. See [78]

for a motivation of why this specific construction of the expansion tensor can yield global information.

To achieve optimal convergence, the prefactor α must be adapted throughout DMRG. The rationale behind this is that the energy can momentarily rise due to the expansion as the weight of Schmidt coefficients and the energy contribution of their corresponding states must not align [78, 83]. This implies that the expansion is typically associated with a slight increase in energy. The prefactor α must hence be tuned down close to convergence. See [78] for a description of the algorithmic procedure on how to best adapt α .

At this stage, no alterations have been made to the state. However, the optimization procedure at site M_{i+1} has the potential to adjust the weights of the states within the expanded leg from zero if these states contribute to energy minimization [78]. This implies that the bond dimension could sustain its increase even post-truncation, resulting in an efficient and accurate DMRG implementation [78, 83]. This particular expansion tensor's application will be further explored in the context of time evolution methods, as detailed in section 2.3.2.

2.3 Time-Dependent Variational Principle (TDVP)

The time-dependent variational principle (TDVP) is not just a tensor network algorithm, but rather a fundamental principle of quantum mechanics [86, 87] that goes back to Dirac [88] and Frenkel [89]. It plays an equivalently fundamental role in time evolution as the Rayleigh-Ritz principle does for the ground state problem. In fact, the Rayleigh-Ritz principle can be obtained from the Dirac-Frenkel principle in the static limit [86]. This foreshadows one of the most significant results in tensor networks, that DMRG can be retrieved as a limiting case of TDVP, namely taking the limit of infinite time steps during imaginary time evolution $\lim_{\delta\tau \rightarrow \infty} e^{-\delta\tau \hat{H}} |\psi\rangle$ [90]. We will refer to the time-dependent variational principle as Dirac-Frenkel principle to avoid confusion with the equally named algorithm.

Given the time-dependent Schrödinger equation

$$i\partial_t |\psi(t)\rangle = \hat{H} |\psi(t)\rangle, \quad (2.39)$$

where \hat{H} is a hermitian operator in the Hilbert space \mathcal{H} and let $\mathcal{M} \subset \mathcal{H}$ be the variational manifold on which the approximation to the wave function $|\psi(t)\rangle$ shall lie. Let $T_{|\varphi\rangle}\mathcal{M}$ denote the tangent space at $|\varphi\rangle \in \mathcal{M}$. The variational principle for the approximate wave function can be obtained by demanding that the time derivative of $|\varphi\rangle$ satisfies

$$\partial_t |\varphi\rangle \in T_{|\varphi(t)\rangle}\mathcal{M} \quad (2.40)$$

at every point and that the error of the approximation is minimal

$$\min_{|\varphi\rangle} \|\epsilon\rangle\| = \partial_t |\varphi\rangle - \frac{1}{i} \hat{H} |\varphi(t)\rangle, \quad (2.41)$$

where $\|\epsilon\rangle\|$ denotes the norm of the error [91]. Demanding the error to be minimal essentially leaves us with an error that is unavoidable due to our approximation, i.e. $|\epsilon\rangle \notin \mathcal{M} \subset \mathcal{H}$. This means that the error $|\epsilon\rangle$ is in the orthogonal complement of the subspace spanned by all states in \mathcal{M} , allowing us to rewrite the above conditions as

$$\partial_t |\varphi\rangle = \hat{P}(|\varphi\rangle) \frac{1}{i} \hat{H} |\varphi\rangle, \quad (2.42)$$

with the orthogonal projection $\hat{P}(|\varphi\rangle) : \mathcal{H} \rightarrow T_{|\varphi\rangle} \mathcal{M}$ given by $\text{Re} \langle \delta\varphi | \hat{P} |\varphi\rangle = \text{Re} \langle \delta\varphi | \varphi\rangle$ for all $|\delta\varphi\rangle \in T_{|\varphi\rangle} \mathcal{M}$ [91]. This allows us to transform the above equation to its more commonly known form

$$\text{Re} \langle \delta\varphi | \epsilon\rangle = \langle \delta\varphi | \partial_t - \frac{1}{i} \hat{H} |\varphi\rangle = 0, \quad \forall |\delta\varphi\rangle \in T_{|\varphi\rangle} \mathcal{M}. \quad (2.43)$$

This results in various advantageous properties such as the conservation of energy and norm during real-time evolution [91] that are not trivially fulfilled by other time evolution algorithms [90, 92]. The Dirac-Frenkel principle is not only foundational to tensor network algorithms, but its near optimality for variational approximation in quantum molecular dynamics simulations has been shown in [91], and it is widely used in other fields, (e.g. the multi-configuration time-dependent Hartree (MCTDH) algorithm used in quantum chemistry [93–95]).

In tensor network language, the projector onto the tangent space of our variational Manifold \mathcal{M} is given as

$$\hat{P}(|\varphi\rangle) = \sum_{i=1}^L P_{i-1} \otimes \mathbb{1} \otimes P_{i+1} - \sum_{i=1}^L P_i \otimes P_{i+1}. \quad (2.44)$$

We note the similarity of this definition to equation (2.41) by recognizing the local nature of our equations. The first term in our projector allows for variations of the site tensor (wave function), and the second term minimizes the error projecting out components parallel to our state by that guarantying norm conservation. Now, similar to equation (2.42), we can insert the projector into our time-dependent Schrödinger equation, which we can formally solve as

$$|\varphi(\Delta t)\rangle = e^{-i\hat{P}(|\varphi\rangle)\hat{H}\Delta t}. \quad (2.45)$$

As the above equation is not yet easily solvable, let us introduce a shorthand notation for a site and bond projected Hamiltonian

$$\hat{P}(|\varphi\rangle)\hat{H} = \sum_i^L H_i^s - \sum_i^L H_i^b \quad (2.46)$$

$$\hat{H}_i^s = P_{i-1} \otimes \mathbb{1} \otimes P_{i+1} \hat{H} \quad (2.47)$$

$$\hat{H}_i^b = P_i \otimes P_{i+1} \hat{H}. \quad (2.48)$$

Applying a symmetric second-order Suzuki-Trotter decomposition to the time evolved state in equation (2.45) gives us [34]:

$$|\varphi(\Delta t)\rangle \approx e^{-i\hat{H}_0^s \frac{\Delta t}{2}} e^{i\hat{H}_0^b \frac{\Delta t}{2}} \dots e^{-i\hat{H}_L^s \Delta t} e^{i\hat{H}_L^b \Delta t} \dots e^{-i\hat{H}_0^s \frac{\Delta t}{2}} e^{i\hat{H}_0^b \frac{\Delta t}{2}} |\varphi\rangle + \mathcal{O}(\Delta t^3). \quad (2.49)$$

We directly identify the DMRG-like sweeping pattern, back and forth through the system. Hence we can equate the time evolution of our state as successively performing two local time evolutions of site and bond tensors: [90]

$$\partial_t M_j = -i\hat{H}_j^{\text{eff}} M_j \quad (2.50)$$

$$\partial_t C_{\underline{j}} = +i\hat{H}_{\underline{j}}^{\text{eff}} C_{\underline{j}}, \quad (2.51)$$

where M_j denotes a site tensor and C_j the center matrix between bond j and $j+1$ denoted by \underline{j} . The extension to the two-site TDVP (2TDVP) is straightforward; one simply has to forward-evolve a two-site tensor $M_{j,j+1}$ and backward-evolve a single-site tensor M_{j+1} , see [90, 92] for a review.

Let us briefly motivate extensions of the standard TDVP algorithm by discussing potential sources of error during time evolution. Typically, four sources of error in TDVP are discussed [92]: The **Trotter error** of order $\mathcal{O}(\Delta t^3)$ per time step; this error can usually be neglected as the prefactor of this scaling is often very small for reasonable bond dimensions [92]. The **Krylov error**, i.e., the error associated with the inexact solution of local equations; this source of error can, at times, be misleading as the absolute error given by the eigensolver (typically Lanczos or Davidson) tends to severely underestimate the actual error persistent in dynamic quantities. The Hochbruck-Lubich criterium often gives a more accurate estimate of the actual error (e.g., the error in Green's functions) [96–98]. The **truncation error** that occurs during the separation of the two-site tensor during 2TDVP; this is typically the main source of error during real-time evolution due to the potentially exponential growth in bond dimension, see chapter 7. Lastly, we have the ominous **projection error**. This error occurs upon the TDSE projection onto our MPS manifold; see equation (2.41). As this error must be orthogonal to the MPS manifold, it immediately follows that it depends upon the bond dimension. No projection error would occur if the manifold would equal the full Hilbert space $\mathcal{M} = \mathcal{H}$. There is no feasible way to measure the projection error during time evolution. However, estimates can be obtained using multi-site variances [99].




Figure 2.11: Local spin-spin susceptibility for the a_{1g} orbital of LiV_2O_4 for various time evolution algorithms (left), see chapter 8 for an in-depth discussion of this material. We find that 2TDVP fails to accurately capture the dynamic of the system, whereas a single time step with a global method, in this case, the global Krylov method (GK), accurately captures the initial spread of information throughout the system and agrees with our subspace expansion methods. The slight differences in (right) can be explained as LSE expands the state more frequently than GSE. Both LSE and GSE have only been used until $\tau = 2$. After that, they follow the same procedure as the curve marked as 2TDVP.

As a rule of thumb, the projection error shrinks upon increasing the bond dimension and the number of active sites involved in the evolution (2TDVP instead of 1TDVP) and grows in the presence of long-range terms in the Hamiltonian. Notably, projection and truncation errors are unaffected by the time step. It can hence, especially for imaginary time evolution ($t \rightarrow i\tau$), be advisable to use larger time steps as the increasing Trotter and Krylov errors are not as substantial, see section 4.3.2 for a discussion about the use of exponentially growing time steps during imaginary time evolution.

The projection error is typically largest in the first time step. The spread of local and global excitations is curtailed if no appropriate basis states encoding the necessary long-range correlations exist. Thus, these excitations are effectively truncated, marking an unrecoverable loss of information. Notable attempts to mitigate this issue have been made, for instance, utilizing global time evolution methods like the global Krylov (GK) solver for the first time step [92] or adapting the TDVP sweeping pattern to mirror the spread of information after a local quench [100]. However, these approaches are not universally applicable due to either their sheer computational cost or their restriction to local quenches. We hence want to direct our focus on subspace expansion techniques for TDVP in the next sections, as they allow for a proper, global, description of quantum dynamics while retaining the efficiency and adaptability of TDVP. See figure 2.11 for a comparison of the aforementioned time evolution methods with subspace expansion techniques.

2.3.1 Global Subspace Expansion (GSE)

The inability of 1TDVP to grow the bond dimension renders it unfit for many applications due to the potentially linear growth of entanglement entropy S during real-time evolution. This, in turn, results in a need for an exponentially growing bond dimension, given by $m \sim e^S$, to accurately describe the system [101–106]. In addition to that, the usual presence of large projection errors at the initial time step severely limits the precision of any TDVP mode that only uses local information, see figure 2.11. Recently, Yang and White [107] proposed a new subspace expansion technique, the global subspace expansion (GSE), that was demonstrated to overcome the aforementioned issues.

We will introduce the algorithm as presented by White and introduce further computational improvements that we developed later on. These improvements were developed in close collaboration with Tizian Blatz, as detailed in [108].

Analogous to the DMRG3S algorithm, the expansion step is local. However, the construction of the expansion tensors is performed globally. GSE employs ancillary states $|\varphi\rangle_n$ to expand the state $|\psi\rangle$ at each bond. While the states $|\varphi\rangle_n$ may be selected arbitrarily, we fundamentally seek to expand the tensor with elements that will facilitate in the next application of the local time evolution operator. Thus, we construct the ancillary states as

$$|\varphi\rangle_n = (1 - \delta t \hat{H})^n |\psi\rangle, \quad (2.52)$$

up to an order k with $n \leq k - 1$. These states can be regarded as an approximation to the expansion of the global time evolution operator to k -th order

$$|\psi(t + \delta t)\rangle = e^{-i\hat{H}\delta t} |\psi(t)\rangle, \quad (2.53)$$

where, unlike usual Krylov subspaces $\mathcal{K}_k(\hat{H}, |\psi\rangle)$ the states are not constructed as $\{\hat{H}^n |\psi\rangle\}_n^k$ but with $\hat{H} \rightarrow (1 - \delta t \hat{H})$. This choice has been proposed in [107] to avoid an additional normalization of ancillary states.

The construction of these ancillary states and the enrichment step precede the actual TDVP update procedure. For a right-to-left sweep the enrichment procedure at site i is given as [107]:

- (i) Perform a singular value decomposition of your original MPS on the leg in sweeping direction

$$M_i \stackrel{\text{SVD}}{=} U_i S_i B_i. \quad (2.54)$$

Note that it is possible to substitute the SVD with a cheaper QR decomposition if one does not wish to truncate at this step.

(ii) Build the projector onto the orthogonal complement of the original state as

$$P_i = \mathbb{1} - B_i^\dagger B_i. \quad (2.55)$$

(iii) Project the active site \tilde{M}_i of all ancillary states $|\varphi\rangle_n$ onto the orthogonal complement of $|\phi\rangle$ perform a direct sum of all projected tensors and perform a singular value decomposition of the result

$$P_i \tilde{M}_i \stackrel{\text{SVD}}{=} \tilde{U}_i \tilde{S}_i \tilde{B}_i. \quad (2.56)$$

If the norm of all site tensors is zero after the projection, jump to step (v). Note that we found it necessary to discard singular values smaller than $\sim 10^{-6}$ to ensure numerical stability. This projection step is the crucial novelty of the expansion method as it ensures the invariance of $|\psi\rangle$ as $\tilde{B}_i B_i^\dagger = 0$, by construction.

(iv) Expand the original right-normalized tensor B_i with its projected ancillary counterpart as

$$B'_i = \begin{bmatrix} B_i \\ \tilde{B}_i \end{bmatrix}, \quad (2.57)$$

to obtain a new right-normalized tensor with an enlarged bond dimension.

(v) Shift the orthogonality center to the left

$$M_{i-1} = A_{i-1} U_i S_i B_i B'_i = A_{i-1} M_i B'_i. \quad (2.58)$$

This ensures that the next site tensor has an appropriately enlarged right leg. The shift of the orthogonality center must likewise be performed for all ancillary states $|\varphi\rangle_n$.

Due to the global nature of this enrichment idea, it is both possible and, in some circumstances, sensible to extend this algorithm to higher-order TDVP schemes. The procedure above can simply be performed before any TDVP update scheme described in section 2.3.

Adaptive GSE

We introduce several enhancements to the foundational GSE algorithm, with the principal objective of diminishing its significant computational overhead. Our first strategy involves integrating the enrichment procedure within the 1TDVP sweep to minimize the frequency of reconstructing environment tensors. This can be achieved by performing the enrichment procedure described above between the forward and backward evolution of 1TDVP. When using symmetric integrators, see equation (2.49), the expansion is performed in the second half time step. Care must be taken as the first expansion happens now after the initial half time step; in order to diminish the projection error, it is hence advisable to perform a standard GSE enrichment step before time evolution.

Furthermore, we propose a heuristic to determine opportune moments for applying the

GSE, leveraging the singular value spectrum of site tensors. Given that singular values are ordered, the magnitude of the smallest singular value in non-saturated bonds serves as a proxy for the truncation error, a metric not directly accessible in the 1TDVP framework. Expansion steps are initiated when singular values exceed a pre-defined threshold; otherwise, the standard 1TDVP is employed. This approach significantly improves GSE's efficiency, considering that the Krylov subspace construction is the most resource-intensive step of the algorithm. This heuristic criterion has demonstrated remarkable stability across various problem sets, as detailed in [108].

Despite these efficiency improvements, GSE is often ill-suited for simulations where computational efficiency is essential. Expansion schemes with more cost-effective construction principles, like the local subspace expansion (LSE) detailed in the next section, typically prevail. Therefore, we recommend using GSE, akin to the Global Krylov method; only for a limited number of initial time steps and subsequently transition to less resource-intensive alternatives.

2.3.2 Local Subspace Expansion (LSE)

The concept of Local Subspace Expansion (LSE) initially surfaced in a preprint [109], yet was omitted in the subsequent peer-reviewed publication [107]. We presented a similar approach in [7], that is significantly more efficient compared to GSE while maintaining accurate results. At times, it even significantly exceeds the precision of the previous state-of-the-art method, 2TDVP, particularly in the context of systems characterized by long-range interactions or constrained bond dimensions [7], see figure 2.11.

This approach employs the expansion tensor originally devised for the DMRG3S algorithm, leveraging its local construction principle that still provides global information [83], given as

$$E_i^k = L_{i-1} E_i^{k-1} W_i, \quad (2.59)$$

with $E_i^0 = M_i$. Note that this expansion tensor does not include a tuneable coefficient α as it is unnecessary in the context of time evolution. It is possible to construct higher-order terms of this expansion tensor, akin to GSE. However, as the enrichment step is performed at least twice as often in symmetric integration schemes, we did not find it necessary to go beyond $k = 1$. However, it is plausible that higher-order expansion tensors combined with fewer overall enrichment steps can further improve the efficiency of this method.

The LSE procedure is, at its core, similar to GSE as it also depends crucially upon a projection onto the orthogonal complement of the current site tensor to ensure the invariance of the state. For a left-to-right sweep, assuming a symmetric integration scheme, the complete LSE procedure at site i is given as [7, 109]:

- (i) Time evolve the active site tensor $M_i(t) \rightarrow M_i(t + \frac{\delta t}{2})$.

(ii) Truncate the active site tensor M_i using a singular value decomposition to obtain

$$M_i = B_i S_i U_i, \quad (2.60)$$

where B_i is a rank three tensor. Unlike GSE, we advise against using a QR decomposition in this step as this constitutes the only singular value decomposition of the site tensor; thus, our only option to discard ineffective basis states.

(iii) Renormalize the singular value tensor after truncation and build the bond tensor $D_i = S_i U_i$. Evaluate heuristically whether to expand or skip the enrichment step.

(a) If no expansion is needed, jump to step (viii).

(b) Else, construct the projector onto the orthogonal complement of the site tensor as $P_i = \mathbb{1} - B_i B_i^\dagger$.

(iv) Construct the expansion tensor up to order k as $E_i^k = L_{i-1} E_i^{k-1} W_i$ and normalize it.

(v) Project the E_i^k onto the orthogonal complement of the site tensor $P_i E_i^k = E_i^\perp$.

(a) If the projection results in an empty expansion tensor, i.e. $\|E_i^\perp\| = 0$; jump to step (viii).

(b) Else, use a singular value decomposition to truncate the projected expansion tensor $E_i^\perp = B_i^\perp S_i^\perp U_i^\perp$. Again, we found a small but finite truncation threshold $\sim 10^{-6}$ necessary for stable, norm-conserving simulations.

(vi) Expand B_i with a direct sum as $\tilde{B}_i = \begin{bmatrix} B_i & B_i^\perp \end{bmatrix}$.

(vii) Expand to bond tensor D_i with a zero tensor to match the size of the expanded bond in \tilde{B}_i . $\tilde{D}_i = \begin{bmatrix} D_i \\ 0 \end{bmatrix}$.

(viii) Calculate new left contractions as $L_i = L_{i-1} \tilde{B}_i^\dagger W_i \tilde{B}_i$. Note that $\tilde{B}_i = B_i$ if the expansion step has been skipped.

(ix) Perform a back evolution of the bond tensor $\tilde{D}_i(t + \frac{\delta t}{2}) \rightarrow \tilde{D}_i(t)$. Again, note that $\tilde{D}_i = D_i$ if the expansion step has been skipped.

(x) Multiply the resulting bond tensor \tilde{D}_i into the next site tensor B_{i+1} .

The efficiency of the Local Subspace Expansion (LSE) method is fundamentally determined by the pace at which the site tensor is expanded. Maintaining a modest, constant maximum bond dimension for the projected expansion tensor, specifically in the range of $m \sim 50 - 100$, consistently produces precise results across various simulations [7].

Chapter 3

Dynamical Mean Field Theory

Electronic structure calculations of strongly correlated materials are notably challenging due to the simultaneous significance of both non-local and local correlations. The existence of incoherent quasiparticles in transition metal oxides over a broad temperature range, attributed to the suppression of the Fermi-liquid scale by Hund's coupling [28], plays a crucial role in the emergence of intricate electronic properties [14–16] which are directly pertinent to experiment and the synthesis of *quantum materials* [17, 18]. Dynamical mean field theory (DMFT) [27] established itself as the most successful numerical approach at addressing these problems, demonstrating good agreement with experiments across a broad spectrum of correlated materials [4, 110–113]. In the following, we will present a short summary of DMFT and provide a detailed discussion of tensor network based impurity solvers and functional materials in subsequent chapters.

Building upon the contributions of Metzner and Vollhardt [114], Georges et al. [115, 116] successfully addressed the Hubbard model in infinite dimensions, presenting a solution for its local Green's function. This laid the groundwork for Dynamical Mean Field Theory (DMFT) [27]. More precisely, Metzner and Vollhardt showed that for infinite dimensions d , the self-energy of the Hubbard model becomes momentum-independent:

$$\lim_{d \rightarrow \infty} \Sigma(\omega, k) \rightarrow \Sigma(\omega). \quad (3.1)$$

It was already shown previously by Kuramoto and Watanabe that the same holds true in the limit of infinite coordination number z^{coord} [117], i.e., an infinite number of nearest neighbors per lattice site. Subsequently, Georges et al. [27] argued that their method for self-consistently calculating the local Green's function of the Hubbard model in infinite dimensions could be extended to systems with finite coordination numbers. They argued that based on the large coordination number in many common 3-dimensional lattice structures, for instance $z^{\text{coord}} = 8$ for a body-centered-cubic lattice (iron, chromium, ...) or $z^{\text{coord}} = 12$ for a face-centered-cubic lattice (gold, silver, ...), they would expect a local self-energy to deviate only slightly from its full solution.




Figure 3.1: Mapping of the lattice model onto a self-consistently determined impurity model. The local interaction Hamiltonian remains the same while the coupling to and within the bath (shown in grey) is determined iteratively.

However, this assumption can break down, especially in the vicinity of phase transitions. Inoue et al. [118] demonstrated that the momentum dependency of the self-energy becomes progressively more important as one approaches the Mott transition from the metallic side. This sparked the creation of many extensions of the DMFT algorithm that aimed to reestablish some momentum dependence of the self-energy. Most notably, we want to mention the dynamic cluster approximation (DCA) [119, 120] that restores momentum dependence by solving the DMFT equations for patches in the Brillouin Zone (BZ) and its real space counterpart cellular DMFT (CDMFT) [110, 112, 121]. Refer to [4, 34] for our work using CDMFT.

For the cause of our derivation, let us assume that the self-energy is indeed local. The DMFT algorithm, as described by Georges and Kotliar [115], introduces an auxiliary impurity problem that can be viewed as a treatment of the lattice system in the atomic limit, i.e., it has the same local interaction Hamiltonian as the full lattice problem. The impurity itself couples to a non-interacting bath, see figure 3.1. The bath and its coupling to the impurity site are described by the non-interacting Green's function $\mathcal{G}(\omega)$, see section 3.1 on how to construct the impurity Hamiltonian. Akin to classical mean-field theory, \mathcal{G} can be identified as a Weiss field. Ultimately, upon self-consistency, the Weiss field \mathcal{G} itself is chosen in such a way that the impurity Green's function G^{imp} mimics the local Green's function G of the lattice problem [34]

$$G^{\text{imp}}(z) = \delta_{ij} G_{ij}(z), \quad (3.2)$$

where z can either take the role of real frequencies ω or Matsubara frequencies $i\omega_n$. As this algorithm is designated to treat translationally invariant problems, i and j are unspecified lattice positions. Initially, \mathcal{G} can simply be constructed by some guess allowing us to solve the impurity problem for G^{imp} and subsequently obtain the self-energy as

$$\Sigma^{\text{imp}}(z) = \mathcal{G}^{-1}(z) - (G^{\text{imp}})^{-1}(z). \quad (3.3)$$

This allows us to construct the lattice Green's function $G(z, \mathbf{k})$ from the Dyson equation by using the locality of the self-energy $\Sigma(z, \mathbf{k}) = \Sigma^{\text{imp}}(z)$ as

$$G^{-1}(z, \mathbf{k}) = G_0^{-1}(z, \mathbf{k}) - \Sigma^{\text{imp}}(z), \quad (3.4)$$

where G_0 is the non-interacting Green's function of the lattice system. We extract the local part of the lattice Green's function by integrating over all momenta \mathbf{k}

$$G^{\text{loc}}(z) := G_{ii}(z) = \left(\frac{1}{2\pi}\right)^d V_d \int_{\text{BZ}} d\mathbf{k} \left[G_0^{-1}(z, \mathbf{k}) - \Sigma^{\text{imp}}(z) \right]^{-1}, \quad (3.5)$$

where we took the continuum limit and the integration is performed in the first Brillouin zone (BZ). V_d denotes the volume of the d -dimensional unit cell [34]. Substituting the local Green's function for the impurity Green's function in equation (3.3) finally yields the DMFT self-consistency condition as

$$(\mathcal{G}^{\text{loc}})^{-1}(z) = (G^{\text{loc}})^{-1}(z) + \Sigma^{\text{imp}}(z). \quad (3.6)$$

We iterate above equations by identifying $\mathcal{G} = \mathcal{G}^{\text{loc}}$ until their deviation decreases below a set convergence threshold $|\mathcal{G}^{\text{loc}} - \mathcal{G}| < \epsilon$ at which point the impurity problem faithfully mimics the lattice problem. A convergence threshold of $\epsilon \sim 10^{-4}$ yields reliable results. We note that the limit of infinite connectivity is not the only limit in which DMFT becomes exact. In the atomic limit, i.e. $t = 0$, all sites can be treated independently. Hence, the self-energy is local as well. Further, DMFT is trivially exact in the non-interacting limit as the self-energy vanishes. To conclude this section, we want to present a step-by-step guide for DMFT.

We start with an initial guess for the Weiss field \mathcal{G} . Typically, one does not choose \mathcal{G} directly but obtains it from equation (3.6) by constructing an initial guess for the self-energy $\Sigma^{\text{imp}}(z)$. DMFT is typically robust against *symmetry-preserving* initial choices of \mathcal{G} or $\Sigma^{\text{imp}}(z)$. However, care must be taken in the vicinity of a phase transition as they often suffer from hysteresis, e.g. the critical interaction strength for the transition from a metallic to a Mott-insulating state typically depends upon whether the system is approached from above or below the transition point, both in terms of interaction strength as well as, but to a lesser extent, initial \mathcal{G} [122]. We emphasize *symmetry-preservation* in this context as DMFT often suffers from a spontaneous emergence of magnetic order that can either be induced by a symmetry-broken \mathcal{G} , e.g. by specifying different Weiss fields for up and down spins, or a computational setup that does not exploit the $SU(2)$ symmetry of the underlying model. Equipped with an appropriately chosen \mathcal{G} , from which we construct an impurity problem, we iterate over the following steps until convergence:

- (i) Solve the impurity problem to obtain $G_{\alpha\beta}^{\text{imp}}(z)$. The indices α and β denote impurity orbitals, see chapter 6 for a detailed discussion of multi-orbital simulations.
- (ii) Construct the self-energy of the impurity model as:

$$\Sigma_{\alpha\beta}^{\text{imp}}(z) = \mathcal{G}_{\alpha\beta}^{-1}(z) - (G_{\alpha\beta}^{\text{imp}})^{-1}(z). \quad (3.7)$$

(iii) Equate the lattice self-energy with the self-energy of the impurity model based on the assumption of locality:

$$\Sigma_{\alpha\beta}(z, \mathbf{k}) = \Sigma_{\alpha\beta}^{\text{imp}}(z). \quad (3.8)$$

(iv) Insert the self-energy into equation (3.4) and perform an integration over the momenta of the lattice Green's function $G^{-1}(\omega, \mathbf{k})$ to obtain the local Green's function

$$G_{\alpha\beta}^{\text{loc}}(z) = \left(\frac{1}{2\pi}\right)^d V_d \int_{\text{BZ}} d\mathbf{k} \left[\left(G_0^{-1}(z, \mathbf{k}) - \Sigma^{\text{imp}}(z) \right)^{-1} \right]_{\alpha\beta}. \quad (3.9)$$

(v) Use the DMFT self-consistency condition to obtain the non-interacting local Green's function $\mathcal{G}_{\alpha\beta}^{\text{loc}}$

$$\left[(\mathcal{G}^{\text{loc}})^{-1} \right]_{\alpha\beta}(z) = \left[(G^{\text{loc}})^{-1}(z) + \Sigma^{\text{imp}}(z) \right]_{\alpha\beta}. \quad (3.10)$$

(vi) Check whether the deviation of $\mathcal{G}_{\alpha\beta}^{\text{loc}}$ from the non-interacting Green's function used to construct the impurity problem $\mathcal{G}_{\alpha\beta}$ falls under a set threshold $|\mathcal{G}_{\alpha\beta}^{\text{loc}} - \mathcal{G}_{\alpha\beta}| < \epsilon$, if convergence is reached exit here. Otherwise set $\mathcal{G}_{\alpha\beta} = \mathcal{G}_{\alpha\beta}^{\text{loc}}$ and construct the impurity problem from it. Jump back to (i).

Typically, convergence is reached after 10-20 iterations. Convergence speed can be improved by mixing, for instance, current and previous \mathcal{G} , see [123]. However, such methods should be employed with care as they might force incorrect results. We recommend using them only if the self-consistency loop oscillates around the correct solution.

3.1 Impurity Models

The basis of all impurity solvers is the construction of the impurity model itself. While different models exist for the simulation of various physical problems like the Hubbard-Kanamori Hamiltonian used to describe transition metal oxides [1, 2, 28] or the periodic Anderson model (PAM) [124, 125] that successfully modeled heavy fermion f-orbital compounds, they usually follow the same simple structure that is already present in the single impurity Anderson model (SIAM) [126]. We, hence, want to introduce the typical 3 part structure of these Hamiltonians and how to subsequently obtain the hybridization function required to construct them on the basis of the SIAM, given by:

$$\hat{H} = \hat{H}_{\text{imp}} + \hat{H}_{\text{bath}} + \hat{H}_{\text{hyb}} \quad (3.11)$$

$$\hat{H}_{\text{imp}} = \mu \sum_{\sigma \in \{\uparrow, \downarrow\}} \hat{d}_{\sigma}^{\dagger} \hat{d}_{\sigma} + U \hat{n}_{\uparrow} \hat{n}_{\downarrow} \quad (3.12)$$

$$\hat{H}_{\text{bath}} = \sum_{l=1}^{L_b} \sum_{\sigma} \xi_l \hat{c}_{l\sigma}^{\dagger} \hat{c}_{l\sigma} \quad (3.13)$$

$$\hat{H}_{\text{hyb}} = \sum_{l=1}^{L_b} \sum_{\sigma} (\gamma_l \hat{d}_{\sigma}^{\dagger} \hat{c}_{l\sigma} + h.c.), \quad (3.14)$$

where \hat{d} and \hat{n} are the annihilation and particle number operators on the impurity site. The creation operators on the L_b bath sites are denoted as $\hat{c}_{l\sigma}^{\dagger}$, and μ embodies all local potentials of the impurity, e.g., the chemical potential or the double counting correction that is relevant in DFT+DMFT simulations. The Hamiltonian of the SIAM is quadratic in the single site operators of the bath $\hat{c}_{l\sigma}^{\dagger}$. They can hence be integrated out to yield [115]:

$$\mathcal{G}(z)^{-1} = z - \mu - \int \frac{d\epsilon}{\pi} \frac{\Delta(\epsilon)}{z - \epsilon} = z - \mu - \Delta(z), \quad (3.15)$$

where $\Delta(\epsilon)$ is given as

$$\Delta(\epsilon) = \pi \sum_l \gamma_l^* \gamma_l \delta(\epsilon - \xi_l). \quad (3.16)$$

Thus we can obtain a hybridization function $\Delta(z)$ for our impurity model that mimics the local problem described by the Weiss field \mathcal{G} as

$$\Delta(z) = z - \mu - \mathcal{G}(z)^{-1}. \quad (3.17)$$

Tensor network based impurity solvers and other ED-like algorithms can only represent a finite number of bath sites due to the exponentially growing Hilbert space dimension. Thus, they must approximate the Weiss field as [27]

$$\mathcal{G}^d(z)^{-1} = z - \mu - \Delta^{\text{discr}}(z), \quad (3.18)$$

where the last term denotes the discretized hybridization function, given as

$$\Delta^{\text{discr}}(z) = \sum_l \frac{\gamma_l^* \gamma_l}{z - \xi_l}. \quad (3.19)$$

A faithful correspondence between the impurity model and the local problem defined by \mathcal{G} can hence be found if $\Delta^{\text{discr}}(z) \approx \Delta(z)$. This can be achieved by appropriate fitting or discretization procedures that will be introduced in chapter 4 and chapter 5.

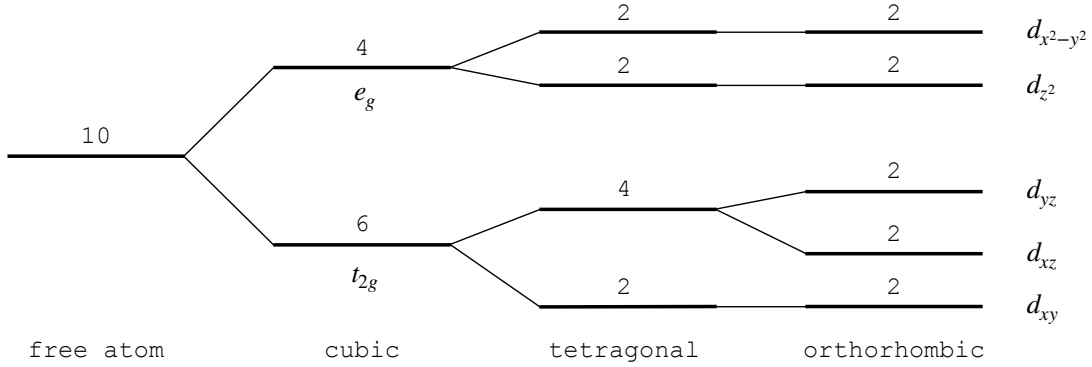


Figure 3.2: Crystal field splitting of d-orbitals under lattice symmetries. Numbers at each level indicate the degeneracy. Adapted from [23].

3.1.1 Real Material Simulations

Materials with partially filled d-or f-shells are typically characterized by strong electronic correlations [19, 27]. Among them are transition metal oxides (TMOs) that are prime candidates for synthetic *quantum materials* due to their myriad of (tunable) functional properties, including superconductivity, Mott transitions, and strange metallicity [14–16].

TMOs are characterized by narrow bandwidths that can be understood as the result of the hybridization of spatially confined transition metal d-orbitals and ligand p-orbitals [23]. However, simple descriptions of such materials are typically unavailable as they are inherently multiband materials; i.e, the Fermi level crosses several bands originating from these hybridized orbitals [28].

Due to the total angular momentum of $L = 2$ in d-orbitals, they have a tenfold degeneracy, one for each of the 5 available L_z quantum numbers and a factor 2 for spin. This degeneracy is lifted in the presence of anisotropic crystal fields that arise in TMOs due to the strong tendency of ligands towards negative valence [23]. Thus, for instance in the presence of a cubic lattice, as found in the TMO SrVO_3 , see [39], the crystal field splits the tenfold degenerate hybridized orbitals into fourfold degenerate, typically energetically less favorable, e_g orbitals and sixfold degenerate t_{2g} orbitals [23], see figure 3.2. Additional degeneracies, especially in partially filled orbitals, can be lifted in the presence of Jahn-Teller distortions [23]. In the context of simulations, this property can lead to crucial simplifications as, depending on the strength of the crystal field splitting, only 3 of the potential 5 orbitals need to be treated in a full DFT+DMFT manner. If the crystal field splitting pushes the e_g bands far enough above the Fermi level, they essentially become uncorrelated. Thus, they can be accurately described by less resource-intensive methods like DFT [20–22] or constrained random phase approximation (cRPA) [127, 128]. We present a comparison of 3 orbital versus full 5 orbital simulations for the heavy fermion TMO LiV_2O_4 in chapter 8.

Nonetheless, any appropriate description of such multiband materials must include inter-orbital couplings, as they constitute the origin of strong correlation in them, especially for materials that are not in the proximity of exhibiting metal-insulator transitions [28]. The appropriate Hamiltonian for the description of an entire d-shell is what is commonly referred to as *full slater* Hamiltonian [129]. In the presence of cubic symmetry, this *full slater* interaction term corresponds to the simpler Hubbard-Kanamori interaction term [28]. Although the Kanamori parametrization is not generally the most appropriate description of TMOs, it is nonetheless commonly used as deviations are expected to be small. A three-band impurity model with Hubbard-Kanamori interaction reads

$$\begin{aligned}
\hat{H}_K = & U \sum_m \hat{n}_{m\uparrow} \hat{n}_{m\downarrow} + U' \sum_{m \neq m'} \hat{n}_{m\uparrow} \hat{n}_{m'\downarrow} + \\
& + (U' - J) \sum_{m < m', \sigma} \hat{n}_{m\sigma} \hat{n}_{m'\sigma} - J \sum_{m \neq m'} \hat{d}_{m\uparrow}^\dagger \hat{d}_{m\downarrow} \hat{d}_{m'\downarrow}^\dagger \hat{d}_{m'\uparrow} \\
& + J \sum_{m \neq m'} \hat{d}_{m\uparrow}^\dagger \hat{d}_{m\downarrow}^\dagger \hat{d}_{m'\downarrow} \hat{d}_{m'\uparrow},
\end{aligned} \tag{3.20}$$

where m, m' run from 1 to 3, U and U' are the intra- and inter-orbital Hubbard interaction, and J is the Hund's coupling [6, 28]. $\hat{d}_{m\sigma}^\dagger$ and $\hat{d}_{m\sigma}$ are fermionic creation and annihilation operators on band m . Under the assumption of spherical symmetry, we have $U' = U - 2J$ [28]. Typically, U' deviates only slightly from $U - 2J$ in non-cubic lattices. Depending on the impurity solver, U' can also be chosen freely. While tensor network based impurity solvers can typically consider arbitrary interactions, other methods, for instance, the numerical renormalization group (NRG), rely heavily on the presence of symmetries. Hence, they often consider yet another simplified version, the Dworin-Narath Hamiltonian [130], which disregards the last term in equation (3.20) (known as pair hopping) and necessarily has $U' = U - J$ [28].

3.2 Impurity Solvers

So far, we glanced over one of the most important aspects of DMFT: the impurity solver. At this point, we solely want to give a broad overview of commonly used impurity solvers; various benchmark results can be found in chapter 7 and chapter 8.

The most commonly used impurity solver is the continuous-time hybridization expansion (CTHYB) [131, 132]. It proved to be highly successful, especially in the context of real material simulations [1, 133–135]. However, it suffers from the presence of a sign-problem, thus effectively restricting it to the Matsubara frequency axis [136–138]. Specifically in the context of large cluster DMFT calculations, the continuous-time interaction expansion algorithm (CTINT) [45] is able to treat a significantly larger number of impurity sites compared to the aforementioned method [139]. However, it is restricted to interaction Hamiltonians that solely consist of density-density terms.

Solver	ω	$i\omega_n$	$T = 0$	T	$\forall \hat{H}^{imp}$	large L_b	large N_i
ED	✓	✓	✓	✓	✓	✗	✗
NRG	✓	✓	✓	✓	✗	✓	✗
CTHYB	✗	✓	✗	✓	✓	✓	✓
CTINT	✗	✓	✗	✓	✗	✓	✓
TN	✓	✓	✓	✓	✓	✓	✓

Table 3.1: **Concise comparison of the characteristics of various impurity solvers.** From left to right, the table summarizes key aspects as follows: Quantum Monte Carlo (QMC)-based methods are generally unable to directly compute Green’s functions on the real frequency axis or at zero temperature. Both NRG and CTINT are restricted in the interaction Hamiltonians they can treat efficiently. The hybridization function requires discretization in ED-like solvers, whereas QMC based approaches typically work directly in the thermodynamic limit. Tensor network based solvers are adept at handling a considerable number of bath sites, L_b , whereas NRG and ED depend on a logarithmic discretization or basis transformations, respectively. Typically, ED and NRG can treat no more than 3 impurity orbitals while results for 5 orbital simulations of real materials are available for CTHYB [140] and TN [39] based solvers. CDMFT results with up to 64 interacting impurities obtained using CTINT are documented in [139].

Note that all implementations are limited to comparably large temperatures due to the $\mathcal{O}(\beta^3)$ scaling of QMC algorithms [132], where β denotes the inverse temperature.

In contrast to this, we want to highlight ED-like solvers that are generally more flexible, allowing them to operate both on the real frequency axis as well as the Matsubara axis for simulations at both finite and zero temperatures. Exact diagonalization was among the first impurity solvers used for DMFT simulations [40]. Nevertheless, the capability to treat a large number of sites is constrained by the exponential growth of the Hilbert space dimension, often restricting ED to rather crude solutions of the problem. Notable effort has been made to significantly enhance the precision of the discretization of the hybridization between the impurity and bath by heavily increasing the number of bath sites L_b and selecting only the most relevant bath sites after a basis transformation [41]. ED is, nonetheless, heavily limited in the number of impurities N_i it can accurately model. Over the years, the Numerical Renormalization Group (NRG) emerged as a benchmark technique due to its exceptional low-energy precision across a broad range of temperatures. However, its exponential scaling severely limits the number of orbitals that can be treated. Moreover, its dependency on non-abelian symmetries and orbital degeneracies restrict it to approximative models that do not capture the full lattice interaction [141, 142].

A major effort of this thesis was developing and improving tensor network impurity solvers. We only want to briefly overview them at this point and provide a detailed discussion in the following chapters. Tensor Network methods have been successfully used for simulations on both the real frequency axis [36, 38, 39, 143–145] and Matsubara frequency axis [1, 2, 4, 33, 35]. Typically, these simulations were restricted to zero temperature. We present a concise scheme to construct finite temperature DMFT simulations with TN network methods and show a comparison with CTHYB data in chapter 8. Additionally, a notable limitation of tensor network solvers on the real frequency axis has been their uniform frequency resolution across the entire spectrum, which typically lags behind the performance of NRG at low frequencies due to the latter’s logarithmic discretization. In chapter 7, we introduce an innovative time evolution technique that allows us to significantly improve the low energy resolution of our methods, enabling us to accurately model the Fermi-liquid behavior of a 3-band model on the real frequency axis. We believe that this advancement paves the way for high-precision simulations of transport properties in transition metal oxides.

Chapter 4

Imaginary Time Impurity Solver

The imaginary time impurity solver (IT) has been a powerhouse of TN-based impurity solvers [1, 2, 4, 33] since its introduction in 2015 by F. Alex Wolf et al.[35], due to its comparably low computational cost. While numerical improvements of the IT solver are the subject of this thesis, its general implementation has been largely performed by Max Bramberger; see [34]. Algorithmic advancements discussed in the following chapters are all contained within our DMFT toolkit *Optimized Basis Tensor Network Impurity Solver* (OTIS), which is a collaborative development with Max Bramberger. All tensor network algorithms employed therein use our highly efficient toolkit SYTEN [77, 78].

The IT solver computes Green's functions directly on the Matsubara axis. Hence, a typically ill-conditioned analytic continuation is required to obtain dynamic quantities that are directly comparable with experiment. See section 4.4 for details on analytic continuation techniques. However, compared to QMC methods, which typically also work directly on the Matsubara axis, it produces noise-free data that can be obtained directly at $T = 0$ or at a wide range of temperatures.

Note that we will exchange z for Matsubara frequencies $i\omega_n$ in our DMFT equations. They are constructed as $\omega_n = \frac{(2n+1)\pi}{\beta^{\text{discr}}}$ where β^{discr} either corresponds to a fictitious temperature β^{eff} that solely sets the energy discretization in our $T = 0$ simulations, or the actual inverse temperature β .

4.1 Discretization

For the sake of generality, let us consider the full matrix-valued hybridization function for multi-orbital impurity models $\Delta_{ij}(i\omega_n)$ where i and j run over all impurity sites. We follow a fitting procedure for the construction of $\Delta_{ij}^{\text{discr}}(i\omega_n)$ that was first introduced by Caffarel and Krauth [40]. We can construct a cost function as:

$$\chi(\gamma_{il}, \xi_{il}) = \frac{1}{N_{\text{max}}} \sum_n^{\text{Nmax}} \omega_n^{-\alpha} \sum_{ij} \left| \Delta_{ij}(i\omega_n) - \sum_l \frac{\gamma_{il}^* \gamma_{jl}}{i\omega_n - \xi_{il}} \right|^2, \quad (4.1)$$

where N_{\max} takes the role of a high energy cut-off in the number of Matsubara frequencies. We found that fitting the hybridization only up to roughly $\omega_n = 6eV$ and disregarding the high energy tail yields better results. As a caveat, we want to stress that all local terms from equation (3.11) must be absorbed into μ , otherwise $\Delta_{ij}(i\omega_n)$ would not decay to zero in the high energy limit, making it impossible to be faithfully fitted with our definition of the cost function [34]. The parameter α can be used to put more emphasis on smaller Matsubara frequencies in the cost function. We found a choice of $\alpha = 0.7$ to yield good results for metallic systems. Hybridization functions of insulating systems are generally easier to fit, and no additional weighting function is required. Hence, we choose $\alpha = 0$ here. As the optimization landscape of the above cost function is, in general, rather complex, we strongly encourage the use of Hessian-based optimization algorithms; see [34] for more details.

The error of our discretization procedure given by the cost function χ is primarily dependent on the bath size L_b and the energy resolution of our simulation, which is determined by β^{discr} (see figure 5.1 for a scaling analysis of the cost function)

4.2 Ground state search

Our objective is to solve the impurity model for the impurity Green's function $G_{ij}(i\omega_n)$. This can be achieved by either computing it directly in frequency space, see [37, 145] or, as it is typically found in modern implementations, obtaining it from performing a Fourier transformation of the time-ordered Green's function in imaginary time τ

$$G_{ij}(\tau) = -\Theta(\tau) \langle \psi_0 | \hat{c}_i(\tau) \hat{c}_j^\dagger | \psi_0 \rangle + \Theta(-\tau) \langle \psi_0 | \hat{c}_j^\dagger \hat{c}_i(\tau) | \psi_0 \rangle. \quad (4.2)$$

Note that we omitted spin indices for the sake of clarity and that $|\psi_0\rangle$ denotes either the ground state of our impurity model for $T = 0$ simulations or a thermofield double state for simulations at finite temperatures. Thus, we must first find $|\psi_0\rangle$ in order to compute the impurity Green's function. The preparation of thermal states will be discussed separately in section 8.3.

Although the ground state search typically does not amount to a large portion of the overall computational cost of a single DMFT iteration, it becomes increasingly more complex for multi-orbital problems. At $\beta^{\text{eff}} = 1000$ the ground state search in LiV_2O_4 , see chapter 8, constituted roughly 8% of the total runtime for a 3-orbital simulation while it already amounts to 38% of the total runtime in the full 5-orbital simulation. This stark contrast in computational cost can be attributed to the rapidly increasing number of candidate quantum number sectors in our ground state search. We only present fundamental concepts in this section, as the ground state search is rather unproblematic for the IT solver compared to the real time solver due to the smaller number of bath sites. More advanced algorithmic developments will be presented in section 5.3.

The quantum number sector of the global ground state is typically not known ahead of time. Therefore, we use DMRG, see section 2.2, in every possible quantum number sector to obtain the ground state of the respective quantum number sector. Let us first clearly define what we mean by quantum number sector. Both single and multi-orbital Hamiltonians conserve well-defined quantum numbers. For instance the SIAM model defined in equation (3.11) conserves the total particle number N ($U(1)$) and the total spin S ($SU(2)$). Hamiltonians of multi-orbital models like the Hubbard-Kanamori Hamiltonian, defined in equation (3.20), also conserve the orbital parity P which yields an additional \mathbb{Z}_2 symmetry.

We can drastically reduce the number of quantum number sectors that we need to consider by constructing a proper initial guess and only considering small deviations around said guess. To obtain a good estimate for the particle number sector N_{guess} of the actual ground state we count the negative (occupied) eigenvalues of our single particle Hamiltonian. We can also restrict the total spin by physical intuition. If we know beforehand that we will, for instance, describe a paramagnetic system, we can safely assume that the total spin will likely be close to $S = 0$. As both of these restrictions are not exact, we typically allow for some variance. We found that a ground state search in $N \in \{N_{\text{guess}} \pm 2N_i\}$ and $S \in \{\pm \frac{N_i}{2}\}$, where N_i denotes the number of impurities, reliably yields the correct quantum number sector. The overall number of quantum number sectors that need to be considered with this approach depends approximately quadratically on the number of impurities. Thus, increasing N_i results in a drastic increase in computational cost.

A possible remedy to this problem is to disregard symmetries of the Hamiltonian altogether. While this approach will likely not result in the ground state, it can drastically limit the number of quantum number sectors that need to be considered by obtaining an even better initial guess for a proper ground state search in individual quantum number sectors.

It can, at times, be beneficial to not exploit the full symmetry of the system for reasons of computational efficiency. This can lead to degenerate ground states. For instance, consider the quantum number sector $\{N = 10, S = 1\}$ where N denotes the particle number and S the total spin. If we disregard the $SU(2)$ of the Hamiltonian and instead only consider a $U(1)$ symmetry conserving the spin S_z , we will find 3 quantum number sectors $\{N = 10, S_z \in \{-1, 0, 1\}\}$ with degenerate ground state energies. The proper, symmetry-conserving, ground state can be reconstructed as an equal-weighted superposition of all its irreducible constituents [34]. Otherwise, an equal-weighted sum cannot be assured as DMRG will typically not construct an equally weighted superposition itself.

4.3 Time Evolution

After acquiring the ground state, we proceed to calculate the time-ordered impurity Green's function as

$$G_{ij}(\tau) = -\Theta(\tau) \langle \psi_0 | \hat{c}_i(\tau) \hat{c}_j^\dagger | \psi_0 \rangle + \Theta(-\tau) \langle \psi_0 | \hat{c}_j^\dagger \hat{c}_i(\tau) | \psi_0 \rangle. \quad (4.3)$$

Again, note that $|\psi_0\rangle$ denotes either the ground state of our impurity model for $T = 0$ simulations or a thermofield double state. The time evolution of the operator $\hat{c}_i(\tau)$ is given as

$$\hat{c}_i(\tau) = e^{\tau(\hat{H}-E_0)} \hat{c}_i e^{-\tau(\hat{H}-E_0)}. \quad (4.4)$$

Here, we add an energy shift $E_0 = \langle \psi_0 | \hat{H} | \psi_0 \rangle$ that serves two purposes: it compensates the factor that we would obtain by applying $e^{\tau\hat{H}}$ to $|\psi_0\rangle$ and, especially relevant for imaginary time-evolution, it ensures numerical stability as the non-unitary time evolution would otherwise quickly reduce the norm below thresholds of numerical accuracy. This leaves us with a simplified definition for our time-ordered Green's function as

$$G_{ij}(\tau) = -\Theta(\tau) \langle \psi_0 | \hat{c}_i e^{-\tau(\hat{H}-E_0)} \hat{c}_j^\dagger | \psi_0 \rangle + \Theta(-\tau) \langle \psi_0 | \hat{c}_j^\dagger e^{\tau(\hat{H}-E_0)} \hat{c}_i | \psi_0 \rangle. \quad (4.5)$$

Thus we have to perform two time-evolutions, one for each excitation \hat{c} (\hat{c}^\dagger). We will refer to the first term in equation (4.5) as particle Green's function and to the second as hole Green's function

$$G_{ij}^p(\tau) = -\Theta(\tau) \langle \psi_0 | \hat{c}_i e^{-\tau(\hat{H}-E_0)} \hat{c}_j^\dagger | \psi_0 \rangle, \quad (4.6)$$

$$G_{ij}^h(\tau) = \Theta(-\tau) \langle \psi_0 | \hat{c}_j^\dagger e^{\tau(\hat{H}-E_0)} \hat{c}_i | \psi_0 \rangle. \quad (4.7)$$

The Θ functions ensure that both time evolutions are performed in positive imaginary time direction (decaying norm, i.e., $e^{-|\tau|(\hat{H}-E_0)}$).

We obtain the impurity Green's function $G(i\omega_n)$ by a Fourier transformation of the time-ordered Green's function. For simulations at zero temperature, we choose to define the Fourier transformation as

$$G_{ij}(i\omega_n) = \int_{-\infty}^{\infty} e^{i\omega_n\tau} G_{ij}^{T=0}(\tau) d\tau. \quad (4.8)$$

This has several advantages that will become apparent once we discuss the usage of linear prediction to extend time series, see section 4.3.1. We use a more commonly found definition for the Fourier transformation for finite temperature Green's functions

$$G_{ij}(i\omega_n) = \int_0^\beta e^{i\omega_n\tau} G_{ij}^T(\tau) d\tau. \quad (4.9)$$

Both of these definitions are equivalent due to the periodicity of Matsubara Green's functions. However, $G_{ij}^{T=0}(\tau)$ and $G_{ij}^T(\tau)$ have to be ordered differently, see figure 4.1.




Figure 4.1: Exemplary depiction of a time-ordered Green's function in our zero temperature ordering (left), see equation (4.8) and the more commonly used definition for finite temperature simulations (right), see equation (4.9).

However, it is significantly easier to accurately account for particle and hole Green's functions that have been evolved, or extended, to different final times τ_{\max} using the definition of equation (4.8). Whenever suitable, we will disregard the predicted part of our time-ordered Green's function and show it in the ordering we use for finite temperature simulations, see figure 4.1 (right).

A subtlety of imaginary time calculations that has not been properly addressed so far is the choice of the final time τ_{\max} . The choice is, arguably, well-defined for calculations at a finite temperature. In this case the time-ordered Green's function is periodic in β , resulting in two time evolutions, each of which is performed until $\tau_{\max} = \frac{\beta}{2}$. Yet, the situation is not as clear-cut for calculations at $T = 0$, i.e., $\beta = \infty$. We consider it good practice to evolve the system in time to at least $T = \frac{\beta^{\text{eff}}}{2}$, or until the Green's function is decayed below a certain threshold. We chose a threshold of 4×10^{-8} . For finite temperature simulations, one should find a smooth transition from particle to hole Green's function at $\tau = \frac{\beta}{2}$. However, it is not obviously apparent why this should hold true for $T = 0$ simulations given their discretization at a fictitious β^{eff} .

We examine this in more detail on the basis of simulations for a 3-orbital Dworin-Narath Hamiltonian with $U=4$ eV and $J=0.6$ eV at half-filling, see section 3.1.1 for a definition of the model. We use a semi-elliptical hybridization function for our calculations defined as

$$-\frac{1}{\pi}\text{Im}\Delta(\omega) = \frac{D}{2\pi}\sqrt{1 - \left(\frac{\omega}{D}\right)^2}, \quad (4.10)$$

for $\omega \in [-D, D]$, where $D = 2$ eV denotes the half-band width in our model. We show the time ordered Green's function in figure 4.2.




Figure 4.2: Zero-temperature time-ordered Green's function for a fully degenerate 3-orbital Dworin-Narath Hamiltonian with $U=4$ eV and $J=0.6$ eV at half-filling. During time evolution, we allowed for up to $m = 4096$ $SU(2)$ states with a time-step of $\delta\tau = 0.05$ and a truncated weight of $t_w = 10^{-11}$. For time evolution we used GSE until $\tau = 5$, 2TDVP until $\tau = 50$ and 1TDVP until the final time of $\tau = \frac{\beta^{\text{eff}}}{2} = 500$. The inset shows the presence of a minute gap between the particle and hole Green's function at $\frac{\beta^{\text{eff}}}{2}$. We chose to present an equal-weighted superposition of the particle and hole component at $\frac{\beta^{\text{eff}}}{2}$. We added the shifted time-ordered Green's function as defined in equation (4.11) as an additional curve in the inset. Even though this gap is of the order of 10^{-6} , it still has a drastic and noticeable effect on the impurity Green's function $G(i\omega_n)$ as shown in figure 4.3.

Calculations were performed at $\beta^{\text{eff}} = 1000$ with a time evolution until $\tau_{\text{max}} = \frac{\beta^{\text{eff}}}{2}$ allowing for up to $m = 4096$ $SU(2)$ states with a time-step of $\delta\tau = 0.05$ and a truncated weight of $t_w = 10^{-11}$. For time evolution, we used GSE until $\tau = 5$, 2TDVP until $\tau = 50$, and 1TDVP for the remainder. We find a small gap between the particle and hole Green's function at $\frac{\beta^{\text{eff}}}{2}$ which, though its size is only in the order of $\sim 10^{-6}$ will lead to unphysical behavior in both the impurity Green's function $G(i\omega_n)$, see figure 4.3 and more importantly the self-energy $\Sigma(i\omega_n)$, see figure 4.4.

To address the issue, we propose two strategies and discuss their effect on both the Green's function and the self-energy of our 3-orbital model. The first method involves the introduction of a shift to the particle/hole Green's functions to close the gap in the time-ordered Green's function. Such a shift can be motivated under the assumption of a constant error during time evolution, for instance, a constant projection error as discussed in section 2.3.

However, it's crucial to note that attempting to directly counteract the underlying causes of these deviations within tensor network algorithms may not be practical due to the resulting exponential increase in computational demands.

We can eliminate the gap at $\frac{\beta^{\text{eff}}}{2}$ by introducing a symmetric shift Δ_{PH} to both the particle and hole Green's function with opposing signs, see inset of figure 4.2:

$$\Delta_{PH} = 0.5(G_{ij}^p(\frac{\beta^{\text{eff}}}{2}) - G_{ij}^h(\frac{\beta^{\text{eff}}}{2})) \quad (4.11)$$

$$G_{ij}^p(\tau) \leftarrow G_{ij}^p(\tau) - \Delta_{PH} \quad (4.12)$$

$$G_{ij}^h(\tau) \leftarrow G_{ij}^h(\tau) + \Delta_{PH}. \quad (4.13)$$

This effectively addresses the oscillations observed in the real part of $G_{ij}(i\omega_n)$. However, closing the gap leaves the imaginary part of the impurity Green's function invariant, see figure 4.3. The persistent oscillations in the imaginary part are indicative of an unphysical solution as the system should adhere to the (generalized) Friedel sum rule, which dictates $\lim_{i\omega_n \rightarrow 0} \text{Im} G(i\omega_n) = 1$. More importantly, while the shift mitigates oscillations in the real part of the self-energy, unphysical behavior in the imaginary part prevails. We find that the expected Fermi-liquid behavior, which anticipates $\lim_{i\omega_n \rightarrow 0} \text{Im} \Sigma(i\omega_n) = 0$, to be completely hidden behind unphysical oscillations. Hence, the gap is not the origin of unphysical behavior in the self-energy, and a simple shift, therefore, does not adequately address the underlying issue. Thus, a second strategy is required using linear prediction. The linear prediction method, as presented in [146], allows us to extend our particle/hole Green's function in τ until they are fully decayed, i.e. smaller than 10^{-10} in magnitude. We want to first discuss the effect linear prediction has on our benchmark system before moving on to a proper introduction of the method in section 4.3.1.

We observe that the real parts of our impurity Green's functions obtained deploying the shift and linear prediction strategies coincide perfectly with each other, see figure 4.3. Moreover, we conclude that the extended decay of the particle/hole Green's functions leads to correct physical behavior in $\text{Im} G(i\omega_n)$ indicated by the accordance with the sum rule. This also translates to the self-energy, where we report a perfect agreement between both strategies in the real part. However, linear prediction additionally uncovers the Fermi liquid behaviour of our system, as shown by the proper decay of the self-energy towards zero, see figure 4.4.

Linear prediction is unavailable in finite temperature simulations as particle/hole Green's functions cannot be extended to arbitrary τ due to their periodicity. Therefore, a standalone discussion of this problem for finite temperature Green's function is warranted and can be found in section 8.3.1.




Figure 4.3: Impurity Green's function $G(i\omega_n)$ for the 3-orbital simulation as defined in figure 4.2. We find that oscillations at low frequencies in the real part of the impurity Green's function are alleviated by the linear prediction method (lp) and the shift method. While we report great mutual agreement of both post-processing methods for $\text{Re}G(i\omega_n)$, the oscillations in $\text{Im}G(i\omega_n)$ remain for the shift due to the periodicity of the Matsubara Green's function. Linear prediction appears to accurately conserve the sum-rule $\lim_{i\omega_n \rightarrow 0} \text{Im}G(i\omega_n) = 1$.




Figure 4.4: Self-energy for the 3-orbital simulation as defined in figure 4.2. We observe large oscillations in $\text{Re}\Sigma(i\omega_n)$ that are mitigated by both the linear prediction (lp) as well as the shift method. We see a clear violation of physics in $\text{Im}\Sigma(i\omega_n)$ by the shift method, while linear prediction captures the Fermi-liquid behavior of the self-energy.




Figure 4.5: Time-ordered Green's function reconstructed via Fourier transformation from an impurity Green's function, which was obtained using linear prediction. Comparison with the pure data reveals an enhanced low energy description as signified by the slower decay in $G(\tau)$, cf text.

Linear prediction mitigates oscillations in our impurity Green's function and effectively restores proper physical behaviour. Yet, we can ask how it deviates from the simple, constant shift method that solely mitigates oscillations. We obtain an answer to this question by reconstructing the time-ordered Green's function, constrained between 0 and β by Fourier transforming the impurity Green's functions that we initially obtained from our extended time ordered Green's function using linear prediction.

The reconstructed time-ordered Green's function shows a significantly larger deviation from the original time-ordered Green's function compared to shift approach, see figure 4.5. Moreover, the reconstructed Green's function does no longer show any gap at $\frac{\beta^{\text{eff}}}{2}$. Considering the deviation from the original time-ordered Green's function, we find that it is, especially for intermediate times, dominated by an exponentially growing term with a nearly constant exponent. However, as the deviation around $\frac{\beta^{\text{eff}}}{2}$ from our original Green's function is large, this exponential term is not simply the result of an insufficiently converged ground state. The overall slower decay, is indicative of a significantly enhanced description of the low energy physics as low lying excited states decay slower under imaginary time evolution. This is consistent with our observation of oscillations at small Matsubara frequencies in our original data. Hence, we conclude that a time evolution until $\frac{\beta^{\text{eff}}}{2}$ does not yield a sufficient description of low energy physics in zero temperature calculations. Moreover, as tensor network based time evolution to a full decay of particle/hole Green's function is unfeasible, time series prediction methods are unavoidable for zero temperature simulations.

4.3.1 Linear Prediction

The derivation presented here is akin to the presentation in [146]. However, as we deem linear prediction essential for zero temperature calculations, we also present it here.

Linear prediction extends a time series $\{x_n\}$ obtained at equidistantly-spaced points $t_n = n\delta t$ in time by fitting it with exponentials of the form $ce^{i\omega\delta t - \eta\delta t}$ with $c \in \mathbb{C}$ and $\omega, \eta \in \mathbb{R}$ allowing it to capture both oscillatory and decaying trends. With the ansatz that new data points can be constructed beyond our latest time step, denoted as t_{obs} , as a linear combination of the p previous data points we obtain

$$\tilde{x}_n = -\sum_{i=1}^p a_i x_{n-i}. \quad (4.14)$$

The coefficients a_i can be determined by minimizing a cost function over a validation interval $(t_{\text{obs}} - t_{\text{fit}}, t_{\text{obs}}]$. Here, t_{fit} should generally be chosen such that it does not include any short time behavior [146, 147]. We can define the cost function χ as

$$\chi(a) = \sum_{n \in N_{\text{fit}}} |\tilde{x}_n - x_n|^2, \quad (4.15)$$

where N_{fit} denotes the number of data points in the validation interval. The minimization procedure can be rephrased as a system of linear equations [146]

$$R\mathbf{a} = -\mathbf{r} \quad (4.16)$$

with

$$R_{ij} = \sum_{n \in N_{\text{fit}}} x_{n-i}^* x_{n-j} \quad (4.17)$$

$$r_i = \sum_{n \in N_{\text{fit}}} x_{n-i}^* x_n. \quad (4.18)$$

Hence, we solve for our coefficients a_i by inverting R . We have not yet shown how our prediction corresponds to exponentials of the form $ce^{i\omega\delta t - \eta\delta t}$. The identification of our prediction with exponentials enables important post-processing steps that ensure the decay of the predicted time series in the long time limit. Let us define the prediction procedure as an iterative application of a matrix A on vectors $X_n := (x_{n-1}, \dots, x_{n-p})^T$ as

$$X_{n+1} = AX_n, \quad (4.19)$$

with

$$A \equiv \begin{pmatrix} -a_1 & -a_2 & -a_3 & \cdots & -a_p \\ 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & \cdots & \cdots & 1 & 0 \end{pmatrix}. \quad (4.20)$$

We can predict our time series m time steps ahead by successive applications of A :

$$x_{n+m} = (X_{n+m})_1 = (A^m X_n)_1. \quad (4.21)$$

This expression can be further simplified by decomposing A into its right eigenvectors as

$$x_{n+m} = (A^m X_n)_1 = (Y D_A^m Y^{-1} X_n)_1 = \sum_{i=1}^p c_i \alpha_i, \quad (4.22)$$

where Y denotes the right eigenvector matrix of A and D_A is a diagonal matrix. The coefficients c_i are given by $c_i = Y_{1i}(Y^{-1} X_n)_i$. The eigenvalues α_i encode resonance frequencies and dampings as $\alpha_i = e^{i\omega_i \delta t - \eta_i \delta t}$ [146]. Since, for our applications, we generally expect a decaying time series, we want to ensure the decay in the long time limit by discarding all eigenvalues α_i with $|\alpha_i| \geq 1$. This can be achieved by simply setting the corresponding coefficient to zero $c_i \rightarrow 0$. Note that discarding a large amount of weight can lead to discontinuities in the time series. To also account for the discarded weight w given as

$$w = \frac{\sum_{i:|\alpha_i| \geq 1} |c_i|}{\sum_{i=1}^p |c_i|}. \quad (4.23)$$

we define a new cost function as

$$\chi(N_{\text{fit}}, N_{\text{obs}}, p) = \max(\epsilon(N_{\text{fit}}, N_{\text{obs}}, p), w(N_{\text{fit}}, N_{\text{obs}}, p)), \quad (4.24)$$

where $N_{\text{fit}/\text{obs}} = \frac{t_{\text{fit}/\text{obs}}}{\delta t}$ and ϵ denotes the maximum difference between our predicted curve and the actually computed time series within our fit interval I_{fit} given as

$$\epsilon = \max_{n \in I_{\text{fit}}} |\tilde{x}_n - x_n|. \quad (4.25)$$

This constitutes an integer-based optimization problem for which we created a custom optimization procedure defined as follows:

- i) Spawn independent processes for N_{runs} random initial configurations $C_i = \{N_{\text{fit}}, N_{\text{obs}}, p\}$.
- ii) Initialize the step size in parameter space Δ_C either as half the integer range between the upper and lower bounds of each parameter or, if no bounds are given, half the initial guess values.
- iii) Spawn N_{walks} process for each configuration C_i with different random seeds.
- iv) Perform a random walk with N_{steps} in parameter space with a step size Δ_C and a randomly drawn number $r \in [-1, 1]$ and map the resulting step in parameter space to an integer $\Delta_{\text{step}} = \text{int}(r \Delta_C)$.

v) Set $\Delta_C \leftarrow \Delta_C/2$. If $\forall \delta_C \in \Delta_C : \delta_C > 1$ jump back to iv) and set the current configuration to the best configuration that has been found during the random walk.

vi) Select the best result from all random walks and initial configurations as your optimization result.

We found that the optimization does not heavily depend upon the initial configuration. Hence, N_{runs} can simply be set to one for most purposes. Therefore, we set N_{walks} to the number of available CPU threads and choose a $N_{\text{steps}} \mathcal{O}(100)$. Note that the computational complexity grows approximately cubically in p . It is thus advisable to provide upper bounds to the optimization procedure. We found an upper bound of $p_{\text{max}} = 200$ to yield reliable results.

Linear prediction is generally a reliable time series prediction method for time-ordered Green's functions in the long-time limit. Moreover, we show that a tensor network based time evolution until $\frac{\beta^{\text{eff}}}{2}$, or even further, is not required as linear prediction manages to accurately predict low-energy physics from significantly earlier times. We applied linear prediction at various times in our previously defined benchmark model and observe excellent agreement between time ordered Green's functions obtained at $\tau = 300 \text{ eV}^{-1}$ and $\tau = \frac{\beta^{\text{eff}}}{2} = 500 \text{ eV}^{-1}$, as shown in figure 4.6. Thus, we can significantly speed up our simulations by limiting the time until which we perform a full tensor network based time evolution.

While we found linear prediction to yield reliable results for functions that are sufficiently well decayed before the onset of the prediction, we nonetheless want to mention notable recent developments in the context of time series prediction. Tian et al. developed a multi-step recursion algorithm in [147] that was shown to produce more accurate predictions if only short time series are available. Lin et al. [148] present a novel approach based on quantic tensor trains. They represent the time series as a MPS and extend the time series based on an entanglement minimization procedure. We want to stress that using a prediction method that could reliably predict $G(\tau)$ from earlier times compared to linear prediction would significantly improve the computational efficiency of our imaginary time solver.




Figure 4.6: Time-ordered Green's functions between $\tau \in [0, \beta]$ (left) obtained from tensor network simulations up to a maximum τ with a further extension using linear prediction. Data is shown for the 3-orbital model defined in figure 4.2. We present the absolute value of the difference between the $\tau = 500 \text{ eV}^{-1}$ curve and curves with shorter maximum evolution times that were extended using linear prediction in the inset. We find that $\tau = 100 \text{ eV}^{-1}$ and $\tau = 200 \text{ eV}^{-1}$ show a stark increase of deviation with increasing prediction time while $\tau = 300 \text{ eV}^{-1}$ and $\tau = 400 \text{ eV}^{-1}$ appear numerically exact. Note that $\tau = 300 \text{ eV}^{-1}$ and $\tau = 400 \text{ eV}^{-1}$ are less accurate at $\tau = 0 \text{ eV}^{-1}$ due to a deviation in the ground state caused by random elements in the discretization and initial state preparation for the ground state search. On the right, we show the corresponding imaginary part of the self-energy. Again, the inset shows the absolute value of the deviation from the $\tau = 500 \text{ eV}^{-1}$ curve. We find that shorter maximum evolution times solely affect the low energy behavior and that the seemingly small deviations in $G(\tau)$ are increased by roughly two orders of magnitude in the self-energy.

4.3.2 How to Perform Efficient Simulations

We want to highlight some of our key improvements in tensor network methods for solving impurity problems on the Matsubara axis, with a focus on optimizing time evolutions, the computationally most intensive part of our simulations. Previously, time evolution strategies combined the global Krylov method with 2TDVP, limiting the former to a single time step due to its high computational cost [34].

Advances in the form of the local and global subspace expansion method now render this combination unnecessary, see figure 2.11 and figure 4.7, significantly boosting efficiency by approximately 2.5 times in zero-temperature simulations of the heavy-fermion compound LiV_2O_4 .

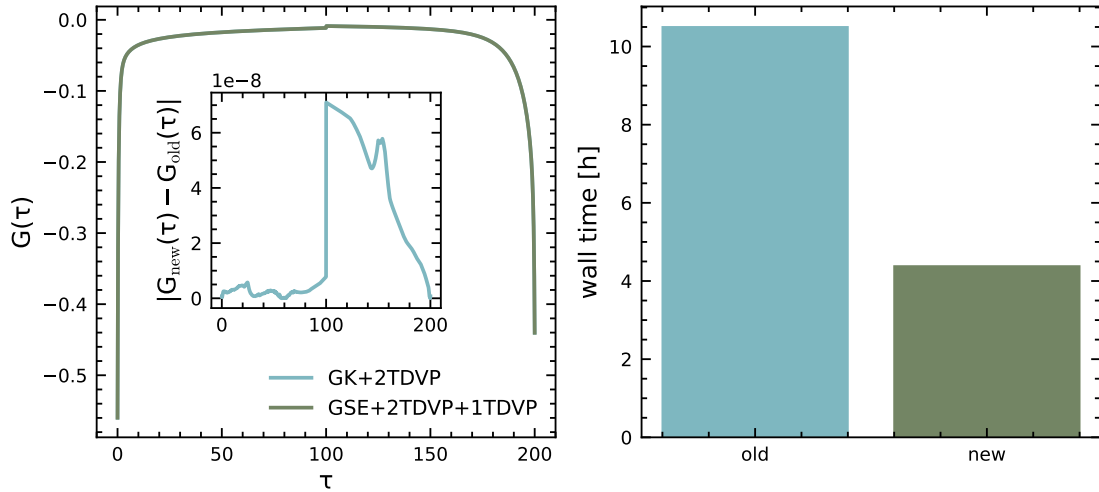


Figure 4.7: Comparison between the old time evolution recipe, a single time step with global Krylov and 2TDVP for the remainder, versus the new time evolution recipe using the GSE method (left). We observe perfect agreement between our time evolutions while achieving a runtime improvement by roughly 2.5. Data is shown for the heavy-fermion compound LiV_2O_4 at $\beta_{\text{eff}} = 200$, see chapter 8. Simulations were performed on 2 Intel(R) Xeon(R) Platinum 8362 processors with 32 CPUs each.

For finite temperature analyses of the same material, these new approaches have demonstrated remarkable time savings, completing tasks in hours that previously exceeded days without performing a single time step. At $\beta = 500$, our new time evolution scheme needs about 11 hours to complete on 4 Intel(R) Xeon(R) Platinum 8362 processors with 32 CPUs each. In comparison the GK did not complete a single time step in the maximally allowed time frame of 5 days. Note that it can sometimes be advisable to use the local subspace expansion method instead of, or in combination with 1TDVP to generate new basis states. However, this is more relevant for real time evolution rather than imaginary time evolution as the latter tends to damp high energy contributions over time.

Selecting the appropriate time step for simulations isn't straightforward, as TDVP errors are not always fully accessible. A larger time step typically reduces overall projection and truncation errors during time evolution but increases errors in the Krylov solver and Trotter scheme, see section 2.3. However, both Krylov and Trotter errors have shown resilience to larger time steps in imaginary time evolution. Possibly because as the state progressively evolves it approaches an eigenstate, thus largely mitigating Krylov errors. Therefore, we want to explore the usage of exponentially growing time steps defined as

$$\delta t_n = \min(\delta t \Lambda^n, \delta t_{\text{max}}). \quad (4.26)$$

Here, Λ determines the growth-rate and δt_{max} marks the largest time step that we allow for to keep the Trotter error in check.




Figure 4.8: Deviation between time-ordered Green’s functions obtained using linear and exponential time steps (left) and the corresponding self-energies (right). The exponential time step is given as $\delta t_n = \min(\delta t \Lambda^n, \delta t_{\max})$ with $\delta t = 0.1$, $\Lambda = 1.01$ and $\delta t_{\max} = 0.5$. The self-energies are indistinguishable from each other; see inset (right). The time-ordered Green’s functions show small deviations from each other, though, as showcased by the self-energy, they are ineffectual.

A choice of $\Lambda = 1.01$ and a capped maximum time step of $\delta t_{\max} = 0.5$ yields good results. Moreover, using equidistant time steps after reaching the cap enables the extension of the Green’s function using linear prediction. We found this approach to yield tremendous efficiency improvements. This results in a cut of runtime for a single iteration of the heavy-fermion compound LiV_2O_4 at $\beta = 2000$ from 31 to 10 hours on the same hardware. While we found exponentially growing time steps to be remarkably stable and accurate, we nonetheless recommend their usage only until convergence is reached. Afterwards, an additional iteration should be performed using only equidistantly spaced time steps to check for any major deviations. See figure 4.8 for a comparison of dynamic quantities computed with exponential time steps.

Further algorithmic improvements will be discussed in chapter 6 when we introduce custom tree tensor network structures for multi-orbital impurity problems.

4.4 Analytic Continuation

Our imaginary time impurity solver offers many computational advantages over its real frequency counterpart, see chapter 5. Typically, real frequency solvers, at best, have an energy resolution of roughly 10 meV for 3-orbital models. Yet, our imaginary time solver accurately resolves strong correlation effects below < 1 meV, see chapter 8. When those results are compared to experimental data much of this precision is lost. For that, Matsub-

ara axis data needs to be mapped to real frequencies. Such a mapping between different complex domains is referred to as analytic continuation.

An analytic continuation between $G(i\omega_n)$ and $G(\omega)$ is given as: [149]

$$G(i\omega_n) = \int_{-\infty}^{\infty} d\omega \frac{\mathcal{A}(\omega)}{i\omega_n - \omega}, \quad (4.27)$$

where $\mathcal{A}(\omega) = -\frac{1}{\pi} \text{Im} G(\omega)$ is the spectral function. The complexity of analytic continuation becomes immediately apparent when we discretize the above equation yielding

$$G(i\omega_n) = \delta\omega \sum_m \frac{\mathcal{A}(\omega_m)}{i\omega_n - \omega_m} = \sum_m K_{nm} \mathcal{A}(\omega_m), \quad (4.28)$$

where $\delta\omega$ denotes the frequency spacing and ω_m the discretized frequencies. The mapping is now given by the analytic continuation Kernel K_{nm} . However, calculating $G(\omega)$ from $G(i\omega_n)$ requires an inverse of K_{nm} . This inversion is generally ill-conditioned, rendering analytic continuation from Matsubara to real frequencies a numerically delicate task. Improvements in terms of stability and accuracy can be achieved by selecting a different Kernel connecting both complex domains. The connection between the time-ordered Green's function $G(\tau)$ and the spectral function \mathcal{A} is given as

$$G(\tau) = - \int_{-\infty}^{\infty} d\omega \frac{e^{-\tau\omega}}{1 + e^{-\omega\beta}} \mathcal{A}(\omega). \quad (4.29)$$

This is typically the Kernel of choice in modern analytic continuation algorithms as it only couples the real parts of $G(\tau)$ and \mathcal{A} , which amounts to more stable optimization procedures [150].

So far, several methods have been devised to stabilize this inversion and enhance the accuracy of analytic continuation. The most widely used among them is the maximum entropy method (MaxEnt) [151–153]. Additionally, the Padé approximation [154] and the recently introduced Nevanlinna algorithm [155] offer alternative approaches, each with its strengths and limitations. However, both of these approaches yield underwhelming results when applied to TN input data. Thus, we will disregard them in our further discussion.

The MaxEnt method constitutes the current benchmark in terms of analytic continuation algorithms. While it can regularly be outperformed by different methods [155] it produces the most reliable results over a wide range of systems and applications. Algorithmically, MaxEnt minimizes a functional of the spectral function given as

$$Q_\alpha[\mathcal{A}] = \frac{1}{2} \chi^2[\mathcal{A}] - \alpha S[\mathcal{A}], \quad (4.30)$$

where χ^2 denotes the difference between the time-ordered Green's function $\tilde{G}(\tau)$ and the current guess $G(\tau) = K\mathcal{A}(\omega)$. Typically, $\tilde{G}(\tau)$ is obtained as an average over multiple

samples of our original time-ordered Green's function with added random noise. The addition of small random noise tends to stabilize the MaxEnt procedure [150]. Let C denote the covariance matrix over the random $G(\tau)$ samples. With this, χ^2 is given as

$$\chi^2[\mathcal{A}] = \sum_{kl} (\tilde{G}_k(\tau) - G_k(\tau))^* C_{kl}^{-1} (\tilde{G}_l(\tau) - G_l(\tau)), \quad (4.31)$$

and S denotes an entropy term given as [152]

$$S[\mathcal{A}] = \int_{-\infty}^{\infty} d\omega \left[\mathcal{A}(\omega) - D(\omega) - \mathcal{A}(\omega) \log \frac{\mathcal{A}(\omega)}{D(\omega)} \right], \quad (4.32)$$

where $D(\omega)$ denotes the default model, a prior knowledge model of the spectral function. Typically, a flat default model is used, which signifies no prior knowledge about the spectral function. Hence, the objective is to deduce the most probable spectral function. The functional Q is minimized across a range of α values, which are usually exponentially distributed. Selecting the optimal α remains difficult as no clearly optimal strategy has emerged. Small α values risk overfitting the noise of the Green's function, while large α values simply make the spectral function resemble the default model $D(\omega)$.

Generally, the optimal α value is determined on the relation of the misfit χ^2 and α . On a log-log scale, this dependence can be characterized by 3 distinct regions [150, 152, 153, 156]. First, the noise-fitting region (small α) where their relation is typically constant. Here, a clear distinction between signal and noise is not always possible, which likely results in incorrect analytic continuations. Second, the transition or information-fitting region in which their relation grows approximately linearly. Typically, optimal values of α are chosen close to or within this region. Finally, the overfitting region (large α) in which their relation is approximately constant, yet again, as typically, only the default model is resuscitated.

Typically, much in the spirit of modern MaxEnt, the goal is to select the optimal α value such that it is as close as possible to the noise-fitting regime, resembling the idea of minimal bias. A robust optimization approach is the line fit procedure [157]. This approach fits lines to the noise and information fitting regions and selects α as the value at their intersection, see figure 4.9.

While MaxEnt constitutes a remarkable improvement from a direct inversion of the Kernel, it remains a bottleneck in the stability and accuracy of imaginary time impurity solvers. MaxEnt primarily suffers from two problems.

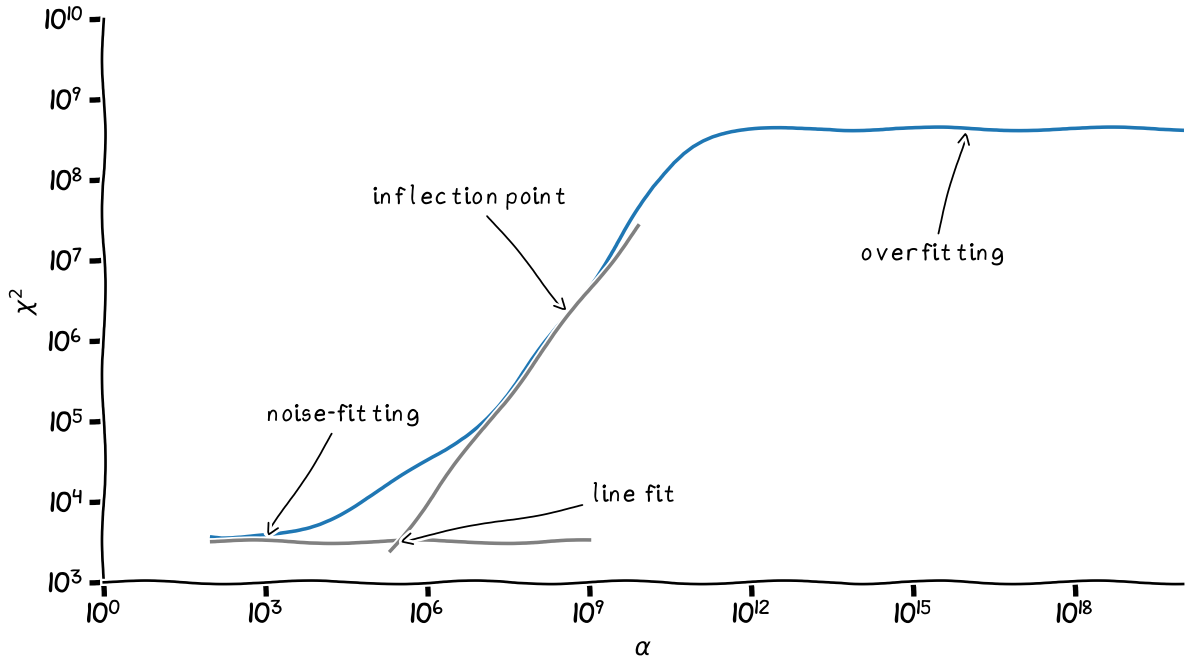


Figure 4.9: Visualization of the dependence of χ^2 on α in an exemplary model and different optimization strategies. MaxEnt does not have any prior knowledge about the spectral function. With prior knowledge, optimization procedures can be greatly stabilized by choosing the value of α at the inflection point. This comes with little to no accuracy penalties when analytically continuing Green’s functions. Yet, we found it critical to use the line fit method for self-energies.

First, due to the nature of imaginary time evolution, high-energy contributions are progressively damped out. Thus, time-ordered Green’s functions tend to have only limited information about high-energy spectral features. Hence, these features succumb to the added random noise, and analytically continued spectral functions cannot accurately capture multiplet structures. The relevance of an accurate high-energy description should not be dismissed. Knowledge of the full range of validity is crucial when comparing our models to experiments. Moreover, in a DMFT context, these comparisons are often based on the self-energy and its derived quantities like the momentum resolved spectral function. Yet, a direct analytic continuation of the self-energy is often unstable [157]. Thus, one typically constructs an auxiliary Green’s function as [158],

$$G_{\text{aux}}(z) = [z - C - \Sigma(z)]^{-1}, \quad (4.33)$$

where the constant C is usually set to $C = \Sigma(i\infty) + \mu$, with μ representing the chemical potential [157]. However, the reconstruction of $\Sigma(\omega)$ from the analytically continued

auxiliary spectral function $\mathcal{A}_{\text{aux}}(\omega)$ requires invoking the Kramers-Kronig relation given as

$$G(z) = \int_{-\infty}^{\infty} d\omega \frac{\mathcal{A}(\omega)}{z - \omega}. \quad (4.34)$$

Here, z simply corresponds to real frequencies with a small broadening, i.e., $z = \omega + i0^+$. Thus, incorrect high-energy structures can influence the accuracy of low-energy features. Ultimately, this limits the accuracy of our imaginary time solver in comparisons to experiment and its predictive power.

The second pitfall of MaxEnt is its strong parameter dependence. As alluded to above, the analytically continued function strongly depends on the correct choice of α . While many suggestions for optimal choices of α have been made [152, 157, 159], none has emerged as clearly superior. Additionally, the result depends strongly on the chosen frequency grid. Finer frequency grids tend to stabilize the optimization procedure and low-energy accuracy. Yet, this is computationally cumbersome as the dimension of the Kernel directly depends on the number of considered frequency points.

In the following section, we present our solution to both of these issues in the form of a new analytic continuation approach, MinKL, that aims to make use of high-quality prior knowledge models. Therefore, it is significantly more accurate compared to current state-of-the-art MaxEnt implementations [150].

4.4.1 MinKL - High Precision Analytic Continuation

Modern MaxEnt algorithms typically focus on the optimization of α and disregard any other parameter dependence. However, the choice of the correct default model can crucially enhance the stability of MaxEnt. Thus, we will explore the role of the default model and showcase its influence on analytically continued functions. This work has been performed in close collaboration with Philipp Westhoff [160].

To begin, let us introduce the Kullback-Leibler (KL) divergence [161] defined as

$$K(P||Q) = \sum_i P_i \log \left(\frac{P_i}{Q_i} \right), \quad (4.35)$$

where P and Q are probability distributions, and P_i and Q_i discrete probabilities defined on the same sample space. It is a relative measure of information that is sometimes referred to as *relative entropy*. It quantifies how much given probability distributions differ from another, where $K(P||Q) = 0$ implies $P = Q$. The entropy $S(P) = -\sum_i P_i \log P_i$ can be obtained as a special case of the KL divergence where Q is a uniform probability distribution. Hence, a maximization of the entropy can generally be equated with a minimization of the KL divergence [161]. The same holds true for the MaxEnt algorithm.

Upon the reasonable assumption that both the spectral function as well as the default model are normalized ($\int_{-\infty}^{\infty} d\omega \mathcal{A}(\omega) = 1$), we can simplify equation (4.32) to

$$S[\mathcal{A}] = - \int_{-\infty}^{\infty} d\omega \mathcal{A}(\omega) \log \frac{\mathcal{A}(\omega)}{D(\omega)}. \quad (4.36)$$

In fact, the lost terms serve no mathematical necessity other than to ensure that $\mathcal{A}(\omega) = D(\omega)$ in the limit of large α [152]. Thus, it is evident that the proper entropic term in MaxEnt simply corresponds to a KL divergence

$$S[\mathcal{A}] = -K(\mathcal{A}||D). \quad (4.37)$$

Hence, we can formulate a functional minimization procedure akin to MaxEnt as

$$Q_{\alpha}[\mathcal{A}] = \frac{1}{2}\chi^2[\mathcal{A}] + \alpha K(\mathcal{A}). \quad (4.38)$$

The definition of χ^2 is still given by equation (4.31) and we define $K(\mathcal{A})$ as

$$K[\mathcal{A}] = \int_{-\infty}^{\infty} d\omega \left[\mathcal{A}(\omega) \log \frac{\mathcal{A}(\omega)}{D(\omega)} - \mathcal{A}(\omega) + D(\omega) \right]. \quad (4.39)$$

Note that we also chose to include the additional terms to stay true to convention. Though, formally, this method simply corresponds to MaxEnt with a different sign, we nonetheless want to refer to it as MinKL as we believe that it reflects more accurately the nature of the algorithm as it puts more emphasis on the default model. While the proper role of the default model in MaxEnt was well understood in older literature [162], the lack of available real frequency results led to an omission of the benefits of a properly chosen default model in recent years.

Tensor network based impurity solvers can operate on the Matsubara axis and the real frequency axis, see chapter 5. Current analytic continuation algorithms limit the accuracy of the imaginary time solver, while real frequency solvers are simply limited by their high computational demand. We propose to alleviate these shortcomings by effectively combining results from both solvers. Real frequency results shall serve as a default model for the MinKL algorithm, thus stabilizing the analytic continuation procedure while preserving the already high low-frequency precision of the imaginary time solver.

We will explore two test systems to demonstrate the method's capabilities. Initially, we will establish MinKL's accuracy using high-precision reference data for the 3-orbital Dworin-Narath Hamiltonian on the real frequency axis that we obtained using our new complex time impurity solver, see chapter 7. Further details of this model can be found in section 4.3 and section 7.4.

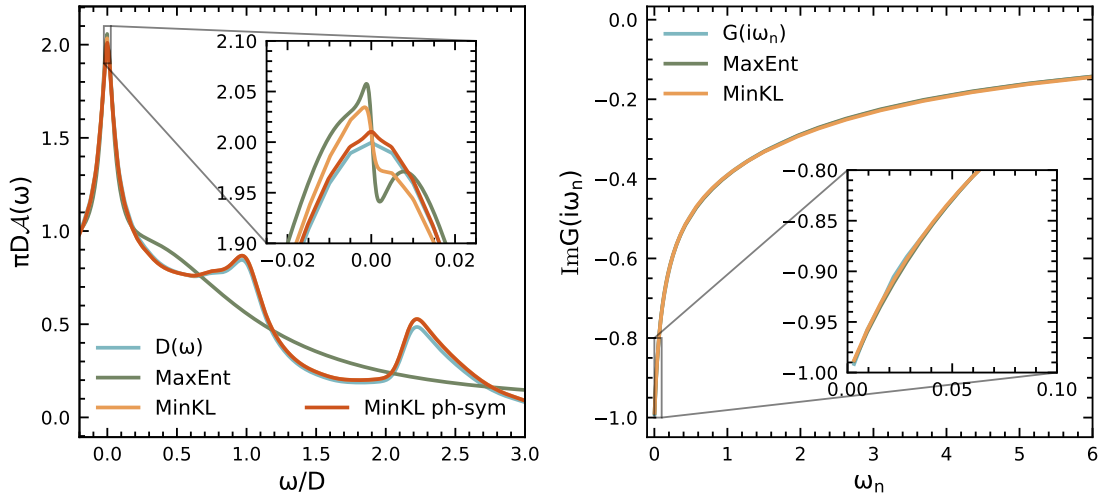


Figure 4.10: Spectral functions for our MinKL default model $D(\omega)$ (left) obtained from our complex time impurity solver, see chapter 7, and analytic continuations of $\beta_{\text{eff}} = 1000$ Matsubara Green's functions using MaxEnt and MinKL. We find overall good agreement between all methods in the low-frequency range. However, MaxEnt fails to accurately capture the multiplet structure of our model and does not show proper decay in the large frequency limit. We notice a particle-hole asymmetry in (left). Hence, we have included a symmetrized version of MinKL. We show the comparison of our input Green's function $G(i\omega_n)$ (right) with the Matsubara Green's functions, which we obtained from our optimal MaxEnt and MinKL spectral functions.

We analytically continue data for the 3-orbital Dworin-Narath Hamiltonian with $U=4$ eV and $J=0.6$ eV at half-filling and $\beta_{\text{eff}} = 1000$. Additional computational parameters are detailed in figure 4.2. We find that the use of our high-precision reference data as a default model significantly improves the analytically continued spectral function. While both MaxEnt and MinKL conserve the (generalized) Friedel sum rule $\pi D_{\mathcal{A}}(0) = 2$, the deviation from our reference data near the Fermi edge is smaller for MinKL. Moreover, MinKL maintains the accurate high-energy description of the default model, while MaxEnt cannot resolve the upper Hubbard bands, see figure 4.10. When we analytically continue our obtained spectral functions back to the Matsubara frequency axis, we find that both MaxEnt and MinKL deviate only slightly from each other, thus corroborating that Matsubara axis data is dominated by low-energy information. All analytically continued data was obtained on a hyperbolic frequency mesh with 4001 points between $\frac{\omega}{D} \in [-4, 4]$.

We have shown that MinKL maintains the accuracy of high-precision reference data. However, the availability of such reference data renders analytic continuation irrelevant. Thus, we want to consider a more realistic scenario, the transition metal oxide LiV_2O_4 , which will be discussed in detail in chapter 8. This material sparked interest due to its emerging heavy-fermion physics at low temperatures [46].




Figure 4.11: Comparison between spectral functions (left) and Matsubara Green’s functions (right) between MinKL and MaxEnt for LiV_2O_4 . We conclude that both methods show a clearly pronounced quasiparticle peak just above the Fermi level. This peak was obscured in our real frequency data used as the default model due to the necessity of a broadening η . MaxEnt data offers a closer resemblance to our input $G(i\omega_n)$ on the Matsubara axis initially. However, this is completely mitigated by our bootstrapping approach.

However, we cannot resolve this behaviour with our real-time solver; the quasiparticle peak near the Fermi edge [163] is hidden due to the employed broadening of $\eta = 0.0075$ eV.

As these complex correlation effects only arise at very low temperatures, they pose severe challenges to current state-of-the-art analytic continuation algorithms [150]. The analytic continuation kernel depends linearly on τ , which in turn grows linearly with the inverse temperature β , see equation (4.29). Hence, the computational cost of MaxEnt and MinKL roughly scale as $\mathcal{O}(\beta^3)$ due to matrix inversions in the optimization procedure. Thus, we make use of the discrete Lehman representation (DLR) [164, 165] that effectively resamples our time-ordered Green’s function in a nearly optimal way [164], compressing it to only fractions of its former size. This leads to a roughly 200-fold decrease in computational cost, which in turn allows us to use significantly finer frequency meshes compared to other implementations, ultimately leading to more accurate results.

This enables us to analytically continue results from our imaginary time impurity solver obtained at an inverse temperature of $\beta = 4000$. We use real frequency data from our real time impurity solver, simulated at zero temperature and $\eta = 0.0075$. We used our new MT3N tensor network, see chapter 6, with 400 bath sites per orbital. The specified broadening required time evolution until a maximum evolution time of $T_{\max} = 120$, which we extended further using linear prediction. We allowed for a maximum bond dimension of $m = 300$ at a truncated weight of $w_t = 10^{-11}$ during time evolution.




Figure 4.12: Parameter dependence of MaxEnt (left) and MinKL (right) for LiV_2O_4 at $\beta = 4000$. We show 120 individual analytic continuations for evenly spaced values of α within the specified parameter region; lines are colored accordingly. MaxEnt results vary greatly with α . For large values of α , individual peaks are no longer discerned, and the pronounced quasiparticle peak at the Fermi edge merges with the side peak. In contrast, MinKL only shows a clear α dependency for frequencies between $\omega \in [-0.05, 0.05]$ eV. The pronounced oscillations next to the quasiparticle peak in MinKL are likely the result of an overfitting of noise, as MinKL is already well into the noise-fitting regime for these values of α . Data was obtained on a hyperbolic frequency mesh with 4000 points between $\omega \in [-4.5, 7.5]$ eV.

We find that MinKL perfectly captures all high-energy peaks present in our default model while maintaining the low-energy information of the Matsubara axis and accurately resolving the quasiparticle peak near the Fermi edge, see figure 4.11. MaxEnt also manages to resolve the quasiparticle peak near the Fermi edge. Yet, it already fails to accurately place the first side peak at around -100 meV. Interestingly, the MaxEnt result shows a closer resemblance to the input Green's function when we compare them on the Matsubara axis. We thus employ an extension to our MinKL scheme, which we will refer to as a *bootstrapping approach*. Given that we have already shown that MinKL preserves the information of highly precise reference data, see figure 4.10, and enhances data from less accurate default models, see figure 4.11, we can posit that we can iteratively use the output of our analytic continuation as a new default model.

Indeed, we find that just two bootstrapping iterations drastically reduce the deviation between the input Green's function and our MinKL result on the Matsubara axis. Additional iterations show diminishing returns. Now, MinKL shows a closer resemblance to

both inputs, $G(i\omega_n)$ from our imaginary time solver and the spectral function from our real time solver. All analytically continued data was obtained using a hyperbolic frequency mesh with 4000 points between $\omega \in [-4.5, 7.5]$ eV. The optimal α was selected as the value at the inflection point, see figure 4.9.

Thus, MinKL accurately combines low-energy information from the Matsubara axis with high-energy information obtained directly on the real frequency axis. Additionally, we want to stress the drastically improved stability of the MinKL approach. The quality of the default model significantly affects the optimization landscape of analytic continuation algorithms. This is evident upon considering the dependence of the results on the optimization parameter α .

We consider LiV_2O_4 at $\beta = 4000$, and select 120 equivally spaced values of α for $\log_{10}(\alpha) \in [4.5, 11.5]$, see figure 4.12. The result of MaxEnt is highly dependent on α , which renders the estimation of correct peak heights and positions difficult. At large values of α , it tends to fuse the quasiparticle peak at the Fermi edge with the side peak at around 10 meV. In contrast, MinKL stably discerns both peaks for all values of α as both of them are already separated in the default model. Note that the 3 distinct optimization regions, see figure 4.9, appear compressed in MinKL. Thus, we are already well within the noise-fitting regime at $\log_{10}(\alpha) = 6$, whereas MaxEnt sits right at the transition between the information and noise-fitting regime. Despite that, MinKL shows a significantly reduced result variability over all considered values of α .

Finally, let us consider the analytic continuation of self-energies. In a DMFT context, most quantities that are comparable to experiment are derived from the self-energy. Yet, direct analytic continuation strategies between Matsubara and real frequency self-energies are highly unstable. While indirect schemes, like the construction of auxiliary Green's functions [158], tend to fare better, analytic continuation of self-energies remains a delicate task.

At $\beta = 4000$, LiV_2O_4 is close to a transition into a heavy Fermi liquid. In a Fermi liquid, we would expect the imaginary part of our self-energy to roughly scale as $\text{Im} \Sigma(\omega) \sim \omega^2$ at lower frequencies [166]. Yet, our real time solver shows no traces of Fermi liquid behaviour, despite simulations being carried out at zero temperature. Interestingly, due to the heavy quasiparticle mass, which is determined by the slope of the self-energy at low energies, even the first point on the Matsubara axis is still far removed from zero.

Upon analytic continuation, MinKL accurately resolves all high-energy peaks in this compound. Moreover, it appears to uncover the emerging Fermi liquid behaviour. At $\omega = 0$, the imaginary part of the self-energy sits at just $\text{Im} \Sigma(\omega = 0) = -0.007$ eV. In comparison, our MaxEnt self-energy sits even higher than our broadened real frequency solver at $\text{Im} \Sigma(\omega = 0) = -0.035$ eV.




Figure 4.13: Analytic Continuation of $\Sigma(i\omega_n)$ for LiV_2O_4 at $\beta = 4000$. Both analytic continuation methods make use of an auxiliary Green's function G_{aux} with a constant $C = 2$. MinKL correctly resolves all high-energy peaks, appearing to even correct for peak heights, in comparison to the broadened default model (left). Moreover, MinKL indicates an onset of Fermi liquid behavior in this compound, which is missed entirely in MaxEnt. This is also reflected on the Matsubara axis (right). MinKL shows a better resemblance with the input data, indicating an overall more stable and accurate analytic continuation. MinKL data is shown after 100 bootstrapping iterations.

Moreover, the validity of our MinKL results is corroborated by a closer resemblance of our self-energy with the imaginary time impurity solver input on the Matsubara axis. Both approaches determined the optimal α via the *linefit* method. Additionally, MinKL results are further enhanced by 100 bootstrapping iterations. Data was obtained on a hyperbolic frequency mesh with 8000 points between $\omega \in [-4.5, 7.5]$ eV. An offset of $C = 2$ was used in the creation of the auxiliary Green's function, see equation (4.33).

Chapter 5

Real Time Impurity Solver

During the early development phase of tensor network impurity solvers, considerable emphasis was put on directly determining the impurity Green's function $G(\omega)$. Prominent approaches include the continued fraction expansion technique by Hallberg et al. [167, 168], and the Dynamic Density Matrix Renormalization Group (DDMRG) most notably investigated in the works of Raas et al. [143, 144] and Nishimoto et al. [169, 170]. Furthermore, significant contributions were made by Chebyshev expansion solvers, pioneered by Wolf et al. [37] and Ganahl et al. [145]. In more recent developments, the focus of tensor network impurity solvers has transitioned to the computation of the retarded impurity Green's function $G(t)$ and its subsequent conversion to $G(\omega)$ via Fourier transformation [171]. Approaches differ in both tensor network structures [6, 36, 39, 172] and time evolution schemes ranging from the Global Krylov (GK) method [35] to the Time Evolving Block Decimation (TEBD) [38, 39] and TDVP in the most recent works [4, 172, 173]. However, computations on the real frequency axis are significantly more difficult. There are primarily two reasons for this. First, the entanglement growth during real time evolution, which we will discuss at length in chapter 7, and second, the necessity for large baths to ensure smooth, physical, self-energies. We will thus present a detailed analysis of various discretization procedures in the next section.

Note that we will exchange z for real frequencies ω in our DMFT equations.

5.1 Discretization

Let us consider the discretized hybridization function given as

$$\Delta^{\text{discr}}(z) = \sum_l^{L_b} \frac{\gamma_l^* \gamma_l}{z - \xi_l}. \quad (5.1)$$

While this constitutes a smooth function on the Matsubara axis, it corresponds to a sum over delta peaks on the real frequency axis. Hence, the choice of the optimal discretization scheme is not as clear as for the IT solver. A comparison purely based on a cost function

is generally not applicable due to the discontinuous nature of equation (5.1) on the real frequency axis. Furthermore, the choice of discretization schemes also affects the entanglement of our tensor network states and the revival time, which indicates the onset of finite size effects during real time evolution.

However, the sum over delta peaks in equation (5.1) will become a sum over Lorentzians upon taking the absolute value, i.e., a smooth function, upon introduction of a broadening $\omega \rightarrow \omega + i\eta$. In this case, we can, similar to our imaginary time solver, define a smooth cost function as

$$\chi(\gamma_{il}, \xi_{il}) = \frac{1}{N_{\max}} \sum_n \sum_{ij} \left| \Delta_{ij}(\omega) - \sum_l \frac{\gamma_{il}^* \gamma_{jl}}{\omega + i\eta - \xi_{il}} \right|^2, \quad (5.2)$$

where we omitted the weight factor compared to equation (4.1) as the real time solver offers equal precision over the full frequency range. Currently, it is unfeasible to solve for a completely unrestricted minimum of this cost function due to the large number of free parameters due to the extensive number of bath sites needed for an accurate representation of the hybridization, see figure 5.1. Hence, we propose to determine the minimum of the cost function as follows. Note that we will assume a diagonal hybridization function in the discussion of our procedure as $\Delta_{ij}(\omega)$ must be diagonalized in the presence of off-diagonal contributions for all discretization procedures presented in this chapter.

i) Initialize outer optimization boundaries I_i^{opt} that will act as a high energy cut-off for the considered frequency range of our discretization procedure.

ii) Separate I_i^{opt} into L_{bath} intervals I_{il} and define the bath on-site potentials ξ_{il} as $\xi_{il} = \frac{I_{il} + I_{il+1}}{2}$.

iii) Perform a least squares optimization as $|\Delta_{ii}(\omega) - \frac{1}{\omega + i\eta - \xi_{il}} |\gamma_{il}|^2|$ to find the optimal hopping terms $|\gamma_{il}|^2$. Evaluate the cost function defined in equation (5.2) on the full frequency range.

iv) Iteratively perform steps i)-iii) within a Nelder-Mead optimization procedure to find the optimal discretization interval I_i^{opt} .

While this approach allows to define a stable, cost function based, optimization procedure, it still has several shortcomings compared to the IT solver. Still, a vast number of bath sites is required, see figure 5.1. Moreover, its viability strongly relies on the presence of a substantial broadening, which ultimately limits the spectral resolution of our real time solver. We will later on refer to this procedure as *fit* method.

So far, in the absence of a broadening, the common wisdom has been to use integration based approaches to discretize the hybridization function.




Figure 5.1: Comparison of the convergence speed of fit-based discretization procedures of the real and imaginary time solver for the e_g^π orbital of LiV_2O_4 , see chapter 8. On the left, we show the cost function defined in equation (5.2) for various broadenings and numbers of bath sites. We find that, while the cost function decreases roughly exponentially with the number of bath sites, the exponent depends heavily on the broadening, rendering the fit discretization unfeasible for small broadenings. In contrast, the discretization procedure on the Matsubara axis (right) converges significantly faster in the number of bath sites. However, the exponent depends upon the inverse temperature, which, akin to the broadening on the real frequency axis, determines the energy resolution in our calculations.

In the following, we provide a brief introduction to these methods, commonly referred to as direct discretization schemes. Similar derivations can be found in [43, 174, 175]. See [176] for a full discussion of matrix-valued discretization procedures. We obtain a discretized approximation to the hybridization function $\Delta(\omega)$ by using the trapezoidal rule as [175]

$$\Delta_{ii}(\omega) = \int_{I_i} \frac{|\gamma_i(\xi_i)|^2}{\omega - \xi_i} \approx \sum_l \frac{|\gamma_i(\xi_{il})|^2 \delta\xi_{il}}{\omega - \xi_{il}}, \quad (5.3)$$

where ξ_{il} are node points with spacing $\delta\xi_{il}$. We retrieve our usual definition of the discretized hybridization function by identifying

$$|\gamma_{il}|^2 = |\gamma_i(\xi_{il})|^2 \delta\xi_{il} \quad (5.4)$$

and the node points ξ_{il} as bath energies [175]. Following Wolf et al. [175], this approximation can be improved by reintroducing the integral formalism as

$$|\gamma_{il}|^2 = \int_{I_{il}} d\xi_i |\gamma_i(\xi_i)|^2 = - \int_{I_{il}} d\omega \frac{1}{\pi} \text{Im} \Delta_{ii}(\omega), \quad (5.5)$$

$$\xi_{il} = \frac{1}{|\gamma_{il}|^2} \int_{I_{il}} d\xi_i \xi_i |\gamma_i(\xi_i)|^2 = \frac{\int_{I_{il}} d\omega \omega \text{Im} \Delta_{ii}(\omega)}{\int_{I_{il}} d\omega \text{Im} \Delta_{ii}(\omega)}, \quad (5.6)$$

where we exchanged the multiplication with an interval $\delta\xi_{il}$ in equation (5.4) with an integration over disjoint intervals I_{il} . The bath energies are computed as weighted averages over said intervals. The same equations, which constitute the heart of direct discretization schemes, can be obtained in a more explicit derivation using a Fourier series expansion that is detailed in [43]. However, both derivations rely on approximations. The quality of the above discretization description clearly depends upon two quantities: the number of intervals and the specific choice of intervals. For instance, NRG relies heavily upon an energy separation in the system. Hence, it requires a logarithmic spacing of the discretization intervals I_l [43]. Tensor network impurity solvers do not suffer from this restriction and can use a wide variety of differently spaced intervals.

The choice of intervals in direct discretization procedures substantially impacts various aspects of our simulations. First and foremost, it affects the accuracy of our discretized hybridization. Moreover, errors in the discretized hybridization function will directly influence the precision of our self-energy, see equation (3.7). Due to an insufficient amount of bath sites or an improperly chosen broadening, strongly peaked discretized hybridization functions can lead to unphysical oscillations in the self-energy that obscure physics. While this effect can be somewhat mitigated by exchanging $\Delta^{\text{discr}}(\omega)$ for $\Delta(\omega)$ in the Dyson equation, as shown for instance in [39], we strongly discourage this approach as it will often, nonetheless, lead to an unphysical behavior in the self-energy. Among the most commonly selected direct discretization schemes are linearly spaced intervals that offer an equal energy resolution over the complete support of the discretization and equally weighted intervals popularized by Wolf et al. [175], which we will denote as *equal* in the following. For the latter, intervals are chosen such that they contain equal fractions of the weight of the bath spectral function $J(\omega) = -\frac{1}{\pi}\text{Im}\Delta(\omega)$. This allows for a better fit in areas of high spectral weight which typically leads to a more accurate description around the Fermi edge. Even better descriptions of the low energy physics can be obtained by the usage of logarithmic discretization strategies. We will present data for both a standard logarithmic grid as well as for a *lin-log* grid that uses logarithmically decreasing intervals near the Fermi-edge and linearly spaced intervals for higher energies.

In the following, we will discuss various discretization schemes on the basis of our two benchmark systems from the prior chapter, the SIAM model and LiV_2O_4 , a transition metal oxide that will be discussed in detail in chapter 8.




Figure 5.2: Accuracy of different discretization schemes, as given by the cost function defined in equation (5.2), for the SIAM model at a broadening of $\eta = 0.01$ (left) and LiV_2O_4 at $\eta = 0.05$ (right). We find that linear discretization intervals perform significantly better for small bandwidths and nearly match the precision of the fit discretization for the SIAM model, while they constitute the worst discretization scheme for real materials. We find that all direct discretization schemes converge to the same cost function value in (right). This is a consequence of the finite broadening. Direct discretization schemes do not account for broadening factors potentially added in the construction of $\Delta^{\text{discr}}(\omega)$. However, as showcased in figure 5.7, direct discretization schemes are generally unable to accurately reproduce hybridization functions with a viable number of bath sites at small or zero broadening.

This distinction in two test cases is certainly of importance as the spectral width of the bath spectral function of the SIAM is typically significantly smaller compared to its real-material counterparts, see figure 5.5 and figure 5.6 for visualization of the hybridization of the SIAM model, a semi-elliptical-density-of-states, and LiV_2O_4 respectively. Overall, the equally weighted intervals offer high accuracy for both test systems, see figure 5.2. While the linear discretization strategy slightly outperforms the *equal* method for the SIAM model, due to the small extension of the bath spectral weight, it constitutes the worst method for LiV_2O_4 . The fit discretization offers the highest accuracy overall. However, note that this comparison is inherently biased towards the fit method as direct discretization schemes are, by construction, unable to reproduce the discretization function at a finite broadening. This is highlighted by the convergence of all intervals to the same cost function value as shown in figure 5.2.




Figure 5.3: Revival times in the SIAM model for equally spaced intervals for various bath sizes (left). The predicted revival times for $L_b = 25$ and $L_b = 50$ are indicated as grey dashed lines. We observe a clear linear dependence of the onset of revivals on the number of bath sites, as indicated by the periodicity in the revivals of $L_b = 25$. The spectral functions (right) stress the importance of understanding when revivals occur ahead of time, as all but one spectral function show clear, unphysical oscillations.

In that sense, we also want to point out a further effect of the discretization procedure: the revival time. Generally speaking, we want to evolve in time for as long as possible in order to improve our spectral resolution according to the Fourier reciprocity. The effect of revivals during time evolution is showcased in figure 5.3, where we find that spectral functions heavily oscillate if we evolve too far in time. The emergence of these revivals can both be easily explained, as well as predicted, if we consider the hybridization function in its Fourier domain representation:

$$\Delta(t) = \int d\omega J(\omega) e^{-i\omega t} \quad (5.7)$$

$$\Delta^{\text{discr}}(t) = \sum_l^{L_b} \gamma_l^* \gamma_l e^{-i\xi_l t}. \quad (5.8)$$

The discretization is now given as a sum over exponentials, and any finite sum over complex exponentials will experience a revival. Luckily, for equally spaced intervals, this revival time can be easily estimated as $t_{\text{max}} = \frac{2\pi}{\delta\xi_l}$, as indicated in figure 5.3. This prediction is strictly speaking only valid in the noninteracting case, which explains the small deviation from the actually observed onset of the revival; however, the basic characteristics of this formula, i.e., the linear dependency of the revival time on the number of bath sites, clearly holds true also in the presence of interaction.




Figure 5.4: Comparison of revival times in the SIAM model with $L_b = 50$ bath sites for different discretization schemes (left). The discretization schemes based on linearly spaced intervals experience the latest revival, while logarithmically spaced schemes are only accurate for short periods. The later revival in the *fit* method compared to the *linear* method can be explained by the truncation of high energy spectral weight during the fitting procedure during the adaption of the frequency boundaries I_i^{opt} . This, however, leads to a shift in the Hubbard side-peaks, as shown in (right).

As indicated before, the occurrence of a revival depends crucially on the energy discretization in our system; however, for non-equidistantly spaced intervals, the prediction and classification of revivals are not as obvious. While the revival occurs significantly earlier for the equal-weight discretization compared to the linearly spaced intervals, the onset of the revival is not as disruptive. This can potentially lead to serious errors during simulations as their onset might go undetected. Overall, the logarithmic discretization schemes experience the earliest revivals, requiring the largest number of bath sites to accurately represent the long-time behavior. We find that the *fit* discretization has a slightly later revival compared to the linear scheme. This is merely a consequence of the fact that the *fit* discretization, per definition, is allowed to disregard high energy spectral weight; hence, the energy discretization $\delta\xi_l$ is finer. This, however, leads to stark deviations in the peak position of the upper and lower Hubbard bands as shown in figure 5.4. We do want to stress that such issues will not be as severe in real-material simulations as they typically do not feature sharp band edges.

While the *fit* discretization displays overall good convergence behavior at finite broadenings, it breaks down for $\eta = 0$, see figure 5.1.




Figure 5.5: Discretized hybridization functions in time (left) and frequency (right) for $L_b = 25$. We find a perfect agreement between our *Fourier-domain* approach with the exact hybridization function and only small deviations for the direct discretization scheme. The main benefit of our approach is based on the idea of obtaining $\Delta^{\text{discr}}(\omega)$ via Fourier transformation. We see a stark deviation in (right) for $\Delta^{\text{discr}}(\omega)$ obtained via Fourier transformation (Fourier-domain) and the direct construction in frequency space using the bath parameters obtained from the *Fourier-domain* method (Fourier-domain ω).

So far, RT solvers were heavily restricted in their frequency resolution due to large entanglement growth during real time evolution. Hence, this did not constitute a severe issue as the limit of $\eta = 0$ was inaccessible. However, with our recent introduction of complex time impurity solvers (CT) [6], see chapter 7, calculations at zero broadening became a possibility and improved discretization procedures a necessity.

Therefore, we introduce a new discretization technique that constructs the discretized hybridization function via Fourier transformation. Thus, we will refer to this new method as *Fourier-domain* scheme. A comparison of the hybridization functions in time defined in equation (5.7) and equation (5.8), shows that both constitute smooth functions even in the absence of a broadening. Hence, it is possible to define a cost-function in time as

$$\chi(\gamma_{il}, \xi_{il}) = \frac{1}{N_{\text{max}}} \sum_n \sum_{ij} \left| \Delta_{ij}(t_n) - \sum_l \gamma_{il}^* \gamma_{jl} e^{-i\xi_{il} t_n} \right|^2 \quad (5.9)$$

to obtain our bath parameters, where $N_{\text{max}} = \frac{t_{\text{max}}}{\delta t}$ denotes the number of time steps up to a specified maximum evolution time. Let us first introduce the basic concepts of this idea based on our two test systems before we give an outline for the actual algorithm. Due to the smoothness of the discretized hybridization function in time, we can *exactly* fit the hybridization function, see figure 5.5.

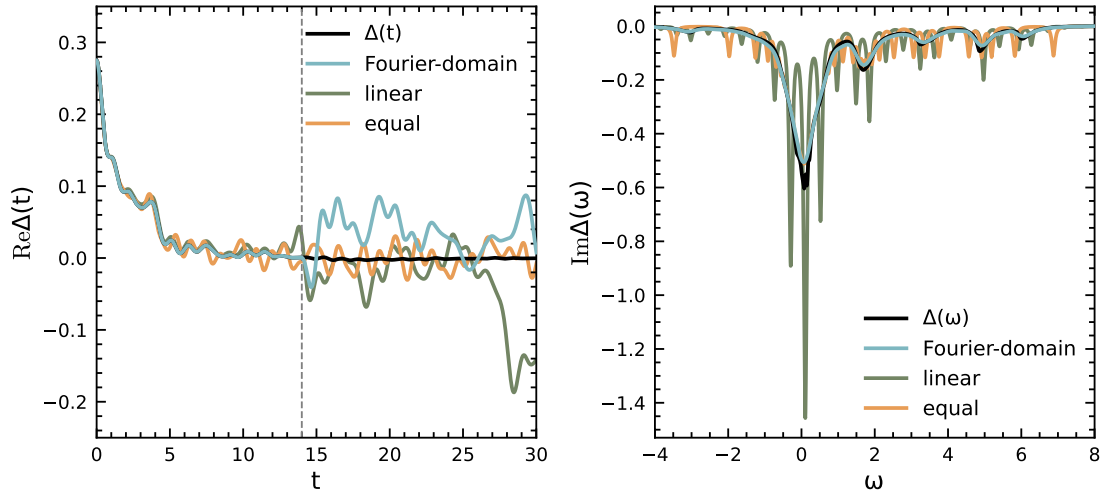


Figure 5.6: Occurrence of revival effects in the hybridization of the e_g^π orbital of LiV_2O_4 for discretizations with $L_b = 50$ bath sites (left). We find a clear breakdown right at the maximally specified optimization time (indicated by the grey line) for the *Fourier-domain* method while it constitutes a numerically exact fit prior. Despite the early onset of a revival, the *Fourier-domain* discretization still significantly outperforms other discretization schemes in the frequency domain (right). All data on the frequency axis is shown for a broadening of $\eta = 0.05$.

While the linear discretization shows only small deviations in the time domain, it deviates significantly in the frequency domain. However, we want to argue that this is largely based on a mathematically ill-posed construction, as shown in figure 5.5. Given that the discretized hybridization is well decayed in time, we can simply construct its frequency correspondent via Fourier transformation, thus avoiding a summation over delta peaks. This leads to a vast improvement in spectral accuracy and allows for an excellent fit of the hybridization, even at a small number of bath sites.

The SIAM model is inherently special as the onset of revivals occurs remarkably late. It is thus necessary to discuss the viability of this approach also for real-materials. The revival in LiV_2O_4 occurs significantly earlier, as shown in figure 5.6. To avoid a complete breakdown of our fitting procedure, it is crucial to know precisely when the revival will occur and limit the optimization to times until the according t_{\max} . We thus construct our algorithm as follows:

- i) Perform a direct discretization with linear intervals.
- ii) Determine the occurrence of the revival as $t_{\max} = \frac{2\pi}{\delta\xi_l}$.
- iii) Construct $\Delta(t) = \int d\omega J(\omega)e^{-i\omega t}$ and minimize equation (5.9), within $t \in [0, t_{\max}]$.

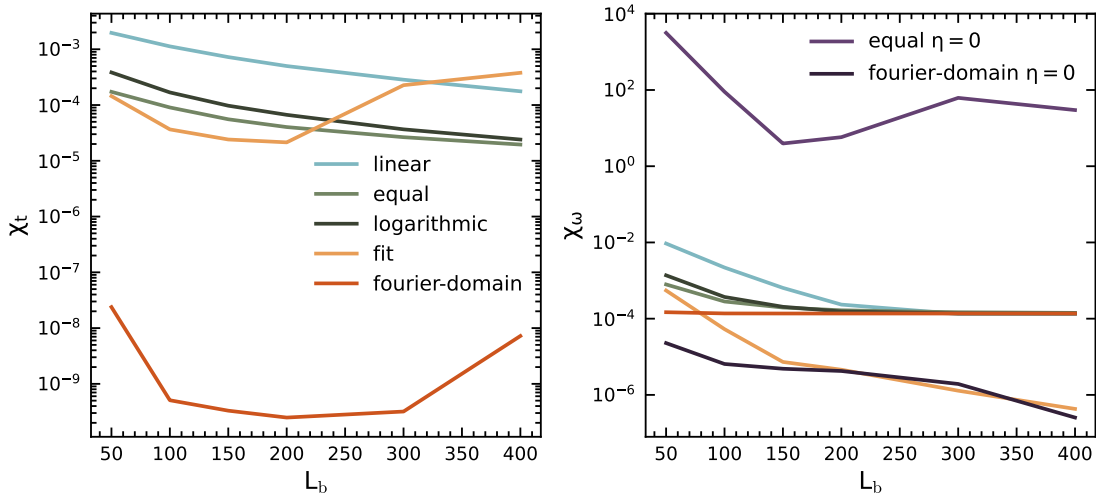


Figure 5.7: Comparison of cost functions for various discretization procedures defined in time (left) and frequency (right) for the e_g^π orbital of LiV_2O_4 . The Fourier-domain scheme shows significant improvements over all other considered discretization procedures for both cost functions. This especially holds true if we consider the added curves in (right) that display a vast improvement in accuracy at zero broadening. If not stated otherwise, data is presented for a broadening of $\eta = 0.05$ on the frequency axis.

iv) Fourier transform $\Delta^{\text{discr}}(t)$ to obtain $\Delta^{\text{discr}}(\omega)$.

This procedure allows us to find stable optima for arbitrary bath sizes and revival times, leaving us with a new discretization procedure that significantly outperforms all other schemes on both the time and frequency axis, see figure 5.7. We report that the *Fourier-domain* discretization outperforms current state-of-the-art discretization procedures with respect to accuracy by about 5 orders of magnitude in the time domain, which in itself is remarkable. However, more importantly, without a broadening, we improve the precision of $\Delta^{\text{discr}}(\omega)$ by up to 8 orders of magnitude compared to the *equal* direct discretization scheme, which was commonly viewed as the most accurate discretization procedure so far [39, 175].

As a curiosity, we want to add that the *Fourier-domain* discretization procedure would open up the possibility of closing the entire DMFT self-consistency loop in time. Such an idea, in a different context, has recently been discussed in [177] where they show that the Dyson equation, in such a scenario, transforms into an integro-differential equation that can be solved with moderate numerical effort. This would completely avoid any numerical instabilities that emerge from the ill-conditioned representation of our quantities of interest on the frequency axis.

5.1.1 Improved Estimators

So far, we found that direct discretization procedures either require a vast number of bath sites or a large broadening in order to generate a smooth discretized hybridization function, see figure 5.7. This is even more pronounced in the context of NRG simulations where, typically, logarithmic grids are used that would necessitate even larger broadenings at higher energies, thus altering the sum rule of spectral functions [178]. In order to avoid this issue, Bulla et al. [178] devised a new strategy for the calculation of the self-energy in impurity models that allows to obtain it without explicitly requiring the discretized hybridization function.

These new strategies are commonly referred to as improved estimators (IE) and can be derived by using first [178] and higher order [179] equations of motion (eom) for impurity Green's function. Using a first-order eom, the self-energy can be computed as [178, 179]

$$\Sigma = FG^{-1}, \quad (5.10)$$

where F denotes an additional correlator given as

$$F(\omega) = -i \int dt e^{i\omega t} \theta(t) \langle \Psi_0 | [V, c_{i,\sigma}(t), c_{i,\sigma}^\dagger] | \Psi_0 \rangle, \quad (5.11)$$

and V denotes the interaction part in our impurity Hamiltonian, which, in its most general form, can be defined as

$$V = \sum_{ijkl} V_{ijkl} c_i^\dagger c_j^\dagger c_k c_l. \quad (5.12)$$

While in theory this amounts to the calculation of additional overlaps during time evolution, this definition conflicts with the possibility of splitting-up the time dependency of the local operator within the correlation function definition, a trick commonly used in real time simulations, see equation (5.38). Thus, improved estimators may necessitate the calculation of one additional time evolution per impurity and spin during real time evolution. However, this is often nonetheless a worthwhile trade-off as improved estimators can drastically increase the accuracy of the self-energy [179].

Now, we can re-express our search for the optimal discretization scheme based solely on the convergence speed of the self-energy in the number of bath sites for various discretization schemes. We consider the SIAM model as described previously. We restricted the maximum time to $t_{\max} = 80D^{-1}$ and the maximum bond dimension to $m = 1536$ during our time evolutions to ensure a meaningful comparison. Let us first establish the convergence speed in the impurity Green's function as this upper bounds our accuracy of the self-energy. We compare various smaller bath sizes to results from a discretization with $L_b = 200$ bath sites. We observe an overall fast convergence of the impurity Green's function in the number of bath sites for all discretization procedures, both in the time and frequency domain. Notably, the equal-weight discretization scheme converges faster than other strategies; see figure 5.8.




Figure 5.8: Comparison of the convergence speed of the retarded impurity Green's function (left) and the impurity Green's function in frequency representation (right) for different discretization procedures. Data is shown for the SIAM model at $U = 2D$, half-filling and $t_{\max} = 80D^{-1}$.




Figure 5.9: Comparison of the convergence speed of the self-energy obtained from the Dyson equation (left) and improved estimators (right) for different discretization procedures. We conclude that the usage of improved estimators significantly boosts the convergence speed of direct discretization procedures. Data is shown for the SIAM model at $U = 2D$, half-filling and $t_{\max} = 80D^{-1}$.

This also holds true for our calculations of the self-energy using improved estimators, see figure 5.9. However, as expected by now, upon using the Dyson equation, the *Fourier-domain* discretization procedure reigns clearly superior.

Thus we may finally conclude our study of discretization procedures. We demonstrated that the *Fourier-domain* discretization procedure clearly outperforms competing strategies. However, it can be bested by other strategies in combination with improved estimators. Nonetheless, we want to stress that *Fourier-domain* discretization reaches the same accuracy with or without improved estimators, thus reinforcing the efficiency and accuracy of our method.

5.2 Optimal Basis

Given that, somewhat independently of the choice of discretization procedure, one needs a large number of sites to accurately represent the impurity model on a tensor network, there has always been a vivid discussion on what the ideal representation might look like. Initially, certainly biased by early tensor network algorithms [143] and NRG [178], a chain-like bath representation, where electronic sites in the bath are connected via nearest-neighbor hopping terms was commonly considered to be ideal [180]. However, Wolf et al. [36] showed that in-fact, a star-like representation, where bath sites are only indirectly connected via second-order hopping processes over the impurity site, has significantly lower entanglement during time-evolution. Several additional methods have been introduced that promise significant performance improvements in impurity solvers based on basis transformations [41, 172, 181–184]. We will thus aim to properly introduce some of the most promising methods below, and we will provide a small comparison study of them based on the mutual information, a quantity often used to measure the entanglement in tensor network states [185, 186].

Let us start our comparison by introducing the two historically most commonly used representations, the star-geometry and the chain-geometry. The star-geometry is the *standard* mapping for impurity models. This is the mapping that we obtain directly from the discretization procedures we discussed in the previous section. The definition of the Hamiltonian in this representation is simply the usual definition for a SIAM model, see equation (3.11):

$$\hat{H} = \hat{H}_{\text{imp}} + \hat{H}_{\text{bath}} + \hat{H}_{\text{hyb}} \quad (5.13a)$$

$$\hat{H}_{\text{imp}} = \mu \sum_{\sigma \in \{\uparrow, \downarrow\}} \hat{d}_{\sigma}^{\dagger} \hat{d}_{\sigma} + U \hat{n}_{\uparrow} \hat{n}_{\downarrow} \quad (5.13b)$$

$$\hat{H}_{\text{bath}} = \sum_{l=1}^{L_b} \sum_{\sigma} \xi_l \hat{c}_{l\sigma}^{\dagger} \hat{c}_{l\sigma} \quad (5.13c)$$

$$\hat{H}_{\text{hyb}} = \sum_{l=1}^{L_b} \sum_{\sigma} (\gamma_l \hat{d}_{\sigma}^{\dagger} \hat{c}_{l\sigma} + h.c.). \quad (5.13d)$$




Figure 5.10: The hopping structure of an impurity model in star-geometry.

Constructing the chain-geometry representation is slightly more difficult and potentially numerically unstable as it involves the Lanczos algorithm that crucially depends on the use of high-precision arithmetic, especially when considering matrix-valued hybridizations [187]. The Hamiltonian in chain-geometry is given as:

$$\hat{H} = \hat{H}_{\text{imp}} + \hat{H}_{\text{pot}} + \hat{H}_{\text{kin}} \quad (5.14a)$$

$$\hat{H}_{\text{pot}} = \sum_{l=1}^{L_b} \sum_{\sigma} \tilde{\xi}_l \hat{c}_{l\sigma}^{\dagger} \hat{c}_{l\sigma} \quad (5.14b)$$

$$\hat{H}_{\text{kin}} = \sum_{\sigma} (\tilde{\gamma}_0 \hat{d}_{\sigma}^{\dagger} \hat{c}_{1\sigma} + h.c.) + \sum_{l=1}^{L_b-1} \sum_{\sigma} (\tilde{\gamma}_l \hat{c}_{l\sigma}^{\dagger} \hat{c}_{l+1\sigma} + h.c.). \quad (5.14c)$$

Here the impurity Hamiltonian \hat{H}_{imp} is unchanged in comparison to the star geometry. The long-range hopping in the star-geometry has now been exchanged for a nearest-neighbor hopping term in \hat{H}_{kin} . This leaves us with a tridiagonal single-particle Hamiltonian that we can construct iteratively using the Lanczos algorithm. We will closely follow [36] in our derivation.

The first term of our Hamiltonian describing the impurity-bath hybridization \hat{H}_{kin} is already fully defined by the initial conditions of the Lanczos procedure

$$\tilde{\gamma}_0 = \sqrt{\sum_l |\gamma_l|^2}, \quad |\tilde{c}_1\rangle = \frac{1}{\tilde{V}_0} \sum_{l=1}^{L_b} V_l |c_l\rangle, \quad (5.15)$$

where $|c\rangle$ and $|\tilde{c}\rangle$ denote single-particle states representing bath sites in star-geometry and chain-geometry respectively. For the Lanczos recursion, it is sufficient to only consider \hat{H}_{bath} , as \hat{H}_{hyb} only has a non-zero contribution in the first element and can thus be ignored otherwise [36].




Figure 5.11: The hopping structure of an impurity model in chain-geometry.

For $n = 2, \dots, L_b - 1$, the Lanczos recursion is defined as:

$$\tilde{\xi}_n = \langle c_n | \hat{H}_{\text{bath}} | c_n \rangle, \quad (5.16a)$$

$$|l_n\rangle = \hat{H}_{\text{bath}} |\tilde{c}_n\rangle - \tilde{\xi}_n |\tilde{c}_n\rangle - \tilde{\gamma}_{n-1} |\tilde{c}_{n-1}\rangle, \quad (5.16b)$$

$$\tilde{\gamma}_n = |\langle l_n | l_n \rangle|^{\frac{1}{2}}, \quad (5.16c)$$

$$|\tilde{c}_{n+1}\rangle = \frac{1}{\tilde{\gamma}_n} |l_n\rangle. \quad (5.16d)$$

Here $n = 1$ constitutes a special case of the recursion, in which case the definition for $|l_n\rangle$ has to be adapted to [36]:

$$|l_1\rangle = \hat{H}_{\text{bath}} |\tilde{c}_1\rangle - \tilde{\xi}_1 |\tilde{c}_1\rangle. \quad (5.17)$$

The coefficients $\tilde{\xi}_n$ and $\tilde{\gamma}_n$ are the newly updated coefficients in chain-geometry representation. Note that the Lanczos algorithm constitutes a unitary transformation between star-geometry and chain-geometry [36]. Hence, the physics must be equivalent in both representations. However, the efficiency of tensor network algorithms heavily depends upon the entanglement in the tensor network state. Thus, we will mainly discuss differences in these mappings on the basis of entanglement measures. We will first discuss the entanglement in the respective ground states on the basis of the mutual information and will later discuss the correlation structure of these representations and the evolution of entanglement with time. Mutual information is closely related to the Kullback-Leibler divergence [188] discussed in section 4.4.1 and tells us how strongly two sites are correlated with each other. We will follow [68, 185, 186] for its introduction in the context of tensor network simulations. For that sake, let us first introduce the von Neumann entropy, defined as

$$S(\hat{\rho}) = -\text{tr}(\hat{\rho} \ln \hat{\rho}) \quad (5.18)$$

where $\hat{\rho}$ is a density matrix. More precisely, we will consider the one- and two-orbital reduced density matrices (RDM) given as

$$\hat{\rho}_i = \text{tr}_{\setminus i} |\psi\rangle\langle\psi|, \quad (5.19)$$

$$\hat{\rho}_{ij} = \text{tr}_{\setminus ij} |\psi\rangle\langle\psi|, \quad (5.20)$$

where the notation $\text{tr}_{\setminus i}$ denotes that we trace out all sites except i . The mutual information is then defined as [68]

$$I_{ij} = (S(\hat{\rho}_i) + S(\hat{\rho}_j) - S(\hat{\rho}_{ij}))(1 - \delta_{ij}). \quad (5.21)$$




Figure 5.12: Comparison between the ground state occupation numbers for the SIAM model at $U = 2D$ and half-filling for star (left) and chain (right) geometry representations. We find a stark deviation in the relative occupations. The on-site energies $\tilde{\xi}$ in chain-geometry show only minor variations from each other leading to a near uniform occupation structure, contrary to the star-geometry where we find a Fermi distribution like occupation.

We will use this concept in a slightly adapted form introduced by Barcza et al. [185] that transforms mutual information into a scalar quantity taking the distance between two sites into account

$$I_{\text{MI}} = \sum_{ij}^L I_{ij} (i - j)^2, \quad (5.22)$$

where L denotes the system size. This measure reflects that tensor network algorithms are mostly local and thus heavily disincentivize long-range correlations in a tensor network state. Applying this measure to a SIAM model at half-filling with $U = 2D$ and $L_b = 20$ leaves us with a mutual information of $I_{\text{MI}}^{\text{star}} = 5.62$ for the star-geometry and $I_{\text{MI}}^{\text{chain}} = 74.48$ for the chain geometry. This discrepancy can, in large parts, be explained by considering the ground state occupation number in both representations. We find, for the SIAM model at half-filling, that sites with a negative on-site energy are nearly fully occupied, see figure 5.12 while sites with a positive on-site energy, whose occupation would constitute an energy penalty, are nearly empty. In all representations, the impurity site has an occupation near 1, which is what we expect for a system at half-filling. However, the chain-geometry produces nearly uniform on-site energies that, in turn, lead to a near-uniform occupation of the ground state. Note that fully empty or fully occupied sites can be represented highly efficiently in tensor network methods as they can be described as product states. We consider the closeness of the star-geometry representation to a Fermi-state like occupation to be the main reason why it posts such a small entanglement compared to the chain geometry.




Figure 5.13: Comparison of the one-particle reduced density matrix of the SIAM model at $U = 2D$ and half-filling with $L_b = 20$ bath sites in star-geometry (left) and chain-geometry (right). Correlations in the star-geometry are nearly perfectly localized. Only the impurity site located in the middle shows small long-range correlations, whereas bath sites in chain-geometry are strongly correlated with each other.

For the SIAM model, we find a roughly 4 times speed-up in the ground state search for the star-geometry, which is very much indicative of the fact that the star-geometry significantly outperforms the chain-geometry in terms of accuracy for the same runtime over a complete DMFT iteration [36]. This is to some extent surprising as the bath representation in chain-geometry features only nearest-neighbor hopping terms, which would lead to vanishing projection errors in two-site update schemes [92, 99]. However, it appears as if the negative impact of the non-local occupation structure of the chain-mapping is far greater than that of the non-local coupling of the star-geometry [36, 189]. We can nicely visualize the localization structure of ground states by considering the one-particle reduced density matrix, see figure 5.13.

The comparison between star- and chain-geometry can be summarized as a comparison between a local correlation structure and a local hopping structure, and clearly, the correlation structure wins. Recently, Kohn et al. proposed a new mapping that would effectively combine the advantages of the chain and the star geometry, a beneficial occupation and localization profile that is similar to the star-geometry while additionally featuring only, at most, next-to-nearest-neighbor hopping terms, see figure 5.14. This is achieved by essentially performing the Lanczos recursion twice, once for the particle bath sites and once for the hole bath sites. Particle/hole bath sites refer to sites whose on-site energies are smaller/larger than the Fermi energy.

They base their mapping, which we will refer to as double-chain mapping, on the thermofield approach [190], a similar mapping has already been presented by de Vega et al. in

the context of open quantum systems [191]. Ultimately, it is based on the same reasoning that we used to explain the efficiency of the star-geometry, the closeness to a Fermi-state. We will closely follow the derivation presented by Kohn et al. [192]. We start by redefining the original single-site operators of our bath as

$$\hat{c}_l \rightarrow \hat{c}_{1,l} \quad (5.23)$$

which allows us to denote their Fermi state as

$$|\text{FS}\rangle = \prod_l \Theta(-\xi_l) \hat{c}_{1,l}^\dagger |\emptyset\rangle, \quad (5.24)$$

where $|\emptyset\rangle$ denotes a vacuum state. We may further introduce a complimentary set of operators acting upon a fictitious ancillary state.

$$\hat{c}_l \rightarrow \hat{c}_{2,l}. \quad (5.25)$$

Note that doubling the Hilbert space is only necessary for finite temperature simulations. This allows us to define a new set of fermionic single site operators as [192]

$$\begin{pmatrix} \hat{f}_{1,l} \\ \hat{f}_{2,l} \end{pmatrix} = \begin{pmatrix} \cos \theta_l & -\sin \theta_l \\ \sin \theta_l & \cos \theta_l \end{pmatrix} \begin{pmatrix} \hat{c}_{1,l} \\ \hat{c}_{2,l}^\dagger \end{pmatrix}. \quad (5.26)$$

This allows us to rewrite \hat{H}_{bath} as

$$\hat{H}_{\text{bath}} = \sum_l \xi_l (\hat{c}_{1,l}^\dagger \hat{c}_{1,l} + \hat{c}_{2,l} \hat{c}_{2,l}^\dagger) = \sum_l \xi_l (\hat{f}_{1,l}^\dagger \hat{f}_{1,l} + \hat{f}_{2,l}^\dagger \hat{f}_{2,l}) \quad (5.27)$$

and \hat{H}_{hyb} as

$$\sum_l \gamma_l \hat{d}^\dagger \hat{c}_{1,l} = \sum_l (\gamma_{1,l} \hat{d}^\dagger \hat{f}_{1,l} + \gamma_{2,l} \hat{d}^\dagger \hat{f}_{2,l}), \quad (5.28)$$

with $\gamma_{1,l} = \gamma_l \cos \theta_l$ and $\gamma_{2,l} = \gamma_l \sin \theta_l$. At $T = 0$, the Fermi distribution becomes a step function. We can thus use $\cos \theta_l = \Theta(\xi_l)$ and $\sin \theta_l = \Theta(-\xi_l)$. Hence, the impurity hybridizes with $\hat{f}_{1,l}$ for $\xi_l > 0$ and $\hat{f}_{2,l}$ for $\xi_l < 0$, effectively separating particle and hole degrees of freedom. At $T = 0$, this allows us to perform a Lanczos tri-diagonalization for each operator species, see equation (5.16). Afterward, we are free to order the resulting particle/hole chains as we please. For the MPS it is advisable to order them as next-to-nearest-neighbor chains [192]. For tree tensor networks one could theoretically spatially separate both chains through the introduction of additional branching nodes, see chapter 6. However, this is counterproductive due to the inefficient scaling of branching nodes.




Figure 5.14: The hopping structure of an impurity model in double-chain mapping.

The full SIAM model, after the double chain mapping, reads:

$$\hat{H} = \hat{H}_{\text{imp}} + \hat{H}_{\text{pot}} + \hat{H}_{\text{kin}} \quad (5.29a)$$

$$\hat{H}_{\text{pot}} = \sum_{l=1}^{L_{\text{particle}}} \sum_{\sigma} \tilde{\xi}_{1,l} \hat{c}_{1,l\sigma}^{\dagger} \hat{c}_{1,l\sigma} + \sum_{l=1}^{L_{\text{hole}}} \sum_{\sigma} \tilde{\xi}_{2,l} \hat{c}_{2,l\sigma}^{\dagger} \hat{c}_{2,l\sigma} \quad (5.29b)$$

$$\hat{H}_{\text{kin}} = \sum_{\sigma} (\tilde{\gamma}_{1,0} \hat{d}_{\sigma}^{\dagger} \hat{c}_{1,1\sigma} + h.c.) + \sum_{l=1}^{L_{\text{particle}}-1} \sum_{\sigma} (\tilde{\gamma}_{1,l} \hat{c}_{1,l\sigma}^{\dagger} \hat{c}_{1,l+1\sigma} + h.c.) \quad (5.29c)$$

$$+ \sum_{\sigma} (\tilde{\gamma}_{2,0} \hat{d}_{\sigma}^{\dagger} \hat{c}_{2,1\sigma} + h.c.) + \sum_{l=1}^{L_{\text{hole}}-1} \sum_{\sigma} (\tilde{\gamma}_{2,l} \hat{c}_{2,l\sigma}^{\dagger} \hat{c}_{2,l+1\sigma} + h.c.). \quad (5.29d)$$

Here the subscript 1 in $\hat{c}_{1,l\sigma}$ denotes the single site operator for the particle sites after the chain mapping, the same notation is used for on-site energies $\tilde{\xi}$ and hopping elements $\tilde{\gamma}$ for particle (1) and hole sites (2) respectively.

Representation	Mutual Information	FS overlap	Runtime [min]
Star	5.621	0.976	2.83
Chain	74.486	0.982	11.71
Double-chain	3.760	0.996	1.82

Table 5.1: Comparison of mutual information and overlaps between a Fermi-state and the actual ground state for different mappings of the SIAM model with $L_b = 20$ bath sites at $U = 2D$ and half-filling. We allowed for a maximum bond dimension of $m = 128$ in our simulations.

Overall this mapping has several advantages. It has a lower mutual information than any of our previously discussed mappings, see table 5.1, while also featuring a beneficial correlation structure. This can, again, be explained by considering the occupation structure in our state, see figure 5.15. We obtain an alternating pattern of nearly fully filled/empty sites marking the separated chains.




Figure 5.15: Occupation structure (left) and the one particle reduced density matrix (right) for the double chain mapping. We report that the state again is very close to being either fully empty or fully occupied while posting a more localized hopping structure that is beneficial for the accuracy in tensor network methods. The RDM shows only minor off-diagonal terms correlating the impurity site in the middle of the chain with the first two sites that show the largest deviation from a Fermi state occupation.

This is also directly reflected by the correlation structure in the one-particle reduced density matrix as we only find significant correlation at the occupied sites and only little correlation between the impurity in the first sites of our new chains. However, the runtime of tensor network impurity solvers is typically dominated by time evolution. Hence, we still have to compare how our representations behave during TDVP.

We find that the double-chain mapping shows both a lower initial entanglement in the excited state and lower entanglement growth throughout time evolution, see figure 5.16. Additionally, we show the entanglement scaling for LiV_2O_4 . Note that LiV_2O_4 is only at quarter-filling. However, the ground state entanglement of the double chain-mapping is, surprisingly, unaffected by that. It is, however, nonetheless plausible that orbitals that show a vast asymmetry in their occupation perform worse in the double-chain mapping compared to the star-geometry. Otherwise, we consider our comparisons as a clear indication of the superiority of the double-chain mapping.

As a concluding remark to this section, we want to briefly discuss further basis-transformations, in particular how we choose to order spin degrees of freedom and basis transformation techniques that lower the entanglement of the state by increasing the bond dimension of the tensor network operator [181, 182, 184, 193].




Figure 5.16: Comparison of the entanglement growth, indicated by the sum over all block entropies, in different representations during time evolution for the SIAM model at half-filling and $U=2D$ for $L_b = 20$ bath sites (left) and for LiV_2O_4 with $L_b = 50$ bath sites (right). Both calculations allowed for a maximum bond-dimension of $m = 1536$. We find that the double-chain mapping is less entangled compared to other representations even throughout time-evolution, thus clearly indicating that it marks a superior mapping.

Typically, spin degrees of freedom in impurity models are degenerate, i.e. they, have an $SU(2)$ symmetry that can, in theory, be exploited in simulations. While this may be necessary to ensure a paramagnetic solution, it is generally computationally more efficient to omit this symmetry and separate spin degrees of freedom from each other, see figure 6.5. This can be achieved by unfolding an electronic site $\{|\emptyset\rangle, |\downarrow\rangle, |\uparrow\rangle, |\downarrow\uparrow\rangle\}$ into two spinless fermionic sites $\{|\emptyset\rangle, |1\rangle\}$. As one can easily see, this does not change the overall Hilbert space dimension of our system but it doubles the number of sites in our tensor network. This introduces a certain additional freedom in the sense that we can now, for multi-orbital impurity models, pose the question of whether it is computationally more efficient to cluster orbital degrees of freedom figure 5.17 or spin degrees of freedom figure 5.18. We find that the deviation in the mutual information for both orderings is within 1% of each other; hence, a simple reordering of bath sites is rather unimpactful. This should come as no surprise as the connectivity of our bath sites is highly limited. Thus, while a simple reordering of bath sites is generally negligible, reordering entire orbitals in multi-orbital real material simulations showed improvements in the mutual information of up to 10% in our simulations.




Figure 5.17: This mapping keeps both spin degrees of freedom close to each other thus minimizing the additional entanglement created through the separation of the interacting sites. The blue/red sites depict sites belonging to a specific orbital, the darker sites depict the impurities.




Figure 5.18: This mapping separates spin degrees of freedom from each other thus trying to spatially separate information. The blue/red sites depict sites belonging to a specific orbital, the darker sites depict the impurities.

Larger entanglement reductions can be obtained when we do not only allow for a reordering of sites but also build superpositions of bath sites. Several attempts at this have been made so far, and while these ideas are often presented as a novelty, their underlying concept is often very similar [41, 182, 184, 193]: the natural orbital basis. The natural orbital basis is a concept from quantum chemistry; more precisely, it is a basis in which the one-particle reduced density matrix of the ground state becomes diagonal [194]. This constitutes a unitary transformation where the transformation matrix can simply be obtained by diagonalizing the one particle reduced density matrix in the computational basis. This approach is, however, not ideal for impurity problems as it would, in general, lead to superpositions between impurity and bath sites. Thus, several custom approaches emerged that tried to construct close approximations to the natural orbital basis that still persevere the locality of the interaction Hamiltonian [41, 182, 193]. While all of these approaches showcased a large reduction of the ground state entanglement, they typically suffer from severe issues during time evolution. The basis does not remain ideal upon excitation and time evolution. This is nicely showcased in [195]. The entanglement will slowly approach the entanglement of the standard computational basis during time evolution. Note that the tensor network operator (TNO) bond dimension is generally enlarged by a natural orbital transformation. Thus, a similar entanglement in the state will lead to an overall decrease in the efficiency of tensor network algorithms. As a consequence, one would need to permanently re-optimize the basis in order to preserve computational efficiency. However, we expect this shortcoming to be eliminated completely with our newly introduced complex time impurity solver, see chapter 7. The reason for this is simple: The imaginary time evolution damps out higher energy contributions to the state. Thus, it will remain close to the lowest-energy eigenstate during time evolution, and it should be possible to construct a lasting, optimal basis for it.

5.3 Ground State Search

While the general approach to the ground state search remains the same in our imaginary time solver, see section 4.2 we are faced with several new challenges. The biggest one is that the computational effort for finding the correct ground state is significantly larger due to the large number of bath sites. While we can still use the same tricks we introduced earlier to reduce the number of candidate quantum number sectors, this is often insufficient to make the ground state search feasible. We will, therefore, introduce two strategies that will significantly reduce the runtime of the ground state search procedure while boosting the overall accuracy of it by constructing better initial states.

We can construct a new, iterative procedure for our ground state search on the basis of our observation that nearly all bath sites in star-geometry or the double-chain mapping are either fully filled or empty. As discussed previously, we already use this assumption to guess the particle number sector; see section 4.2. However, this typically slightly deviates from the correct particle number sector for intermediate coupling strengths, and more importantly, it does not yet restrict other symmetries of our Hamiltonian, e.g. orbital parity. We can obtain a cheap and typically very accurate guess for the complete quantum number sector by performing what we will dub *iterative* ground state search. The idea is to iteratively add bath sites that are energetically closest to the Fermi edge to our impurity cluster until the quantum number sector only changes by double occupations for a predefined number of consecutive iterations. The procedure can be defined as follows:

- i) Construct TNO and TNS for the impurity model without any bath sites and find the correct ground state quantum number sector.
- ii) Count the number of consecutive iterations n_{count} in which the quantum number sector of the ground state changed by simply adding two particles per impurity.
- iii) Add the closest remaining bath sites below and above to Fermi edge for each impurity to the current system and find the new ground state quantum number sector. If the quantum number sector only changed by an added double occupation per impurity, increase n_{count} by 1, otherwise set $n_{\text{count}} = 0$.
- iv) Repeat iii) until the quantum number sector only changes by two particles per impurity for a given number of iterations $n_{\text{count}} > n_{\text{break}}$.
- v) Calculate the quantum number sector for the full system as $N, S = N_{\text{iter}} + 2 * N_{\text{remaining}}, S_{\text{iter}}$ where $N_{\text{remaining}}$ denotes the number of remaining bath sites in our system with negative on-site energies.




Figure 5.19: Runtime comparison for a multi-orbital impurity model with $L_b = 100$ bath sites per orbital at half-filling (left) between our iterative ground state search procedure and a full ground state search around a guessed particle number sector. In (right) we show the significance of orbital degeneracies on the number of quantum number sectors that need to be considered. Note that the comparison in (left) has been performed for a fully degenerate system.

This procedure allows us to obtain the full quantum number sector as double occupancies do not alter the spin symmetry or even the $Z(2)$ symmetry originating from the orbital parity in the Hubbard-Kanamori Hamiltonian. Additionally, we want to stress that it even manages to correctly identify degenerate quantum number sectors, as these degeneracies are also present during the iterative procedure. Care must thus be taken to still increase n_{count} if the ground state quantum number sector is one of the degenerate ground states but not exactly the same as in the prior iteration. It is beneficial to use the full symmetry of the Hamiltonian in this step and, if need be, to change the symmetries under which the tensor network transforms as an irreducible representation once the correct quantum number sector has been identified. We found that already $n_{\text{break}} = 3$ allows to reliably find the correct quantum number sector. We showcase the potential of our new method based on a multi-orbital impurity model at half-filling with a semi-elliptical density-of-states and Hubbard-Kanamori interaction with $U = 2D$ and $J = 0.3D$. We find a significant speed up for our iterative ground state search procedure over the full ground state search around a guessed quantum number sector, see figure 5.19. Note that this large speed-up was obtained for a fully-degenerate system. Given the exponential growth of particle number sectors that need to be considered in the absence of orbital symmetries, the relative advantage of the iterative ground state search method can be expected to show a similar exponential growth. As the final DMRG in the full system typically dominates the overall runtime of the procedure, it does not critically depend upon the presence of degeneracies.




Figure 5.20: Overlaps between the ground state and different initial states in the SIAM model at half filling and $U = 2D$ (left) and the impact of the interaction strength on the overlap with a Fermi state (right).

So far, we focused on reducing the number of different DMRG instances that we have to spawn simultaneously; now, let's consider how we can speed up the convergence of each DMRG simulation. Given that typical RT DMFT simulations consist of $\mathcal{O}(1000)$ sites, the choice of a good initial state becomes crucial. Otherwise, one must allow for excessive sweeping to produce the correct ground state occupation or might be stuck in a local minima altogether. Hence, similar to our argument in section 4.2, we exploit that our target state will be close to a Fermi state in the star geometry or the double-chain mapping. We can construct our Fermi state initial guess as follows:

- i) Diagonalize the single particle terms of the Hamiltonian and obtain eigenvalues D and eigenvectors U .
- ii) Construct N superpositions of creation operators as $\hat{c}_k^\dagger = \sum_l U_{lk} \hat{c}_l^\dagger$ where N denotes the number of particles and k runs over the N sites with smallest eigenvalues.
- iii) Apply the resulting operators to a vacuum state. Note that if no exact application is feasible, a variational application of \hat{c}_k^\dagger with truncation is advisable [92, 196].

To benchmark the effectiveness of this approach, we will consider the SIAM model at half-filling for various interaction strengths, figure 5.20. We compare our results with a *complete* state and what we will refer to as poor man's Fermi state, i.e., a state that can be constructed by setting $U = \mathbb{1}$. The *complete* state is the default construction principle for initial states in SYTEN. It randomly applies all single site operators until it creates every possible basis on every bond [77]. We strongly discourage this construction principle

as it is incredibly slow and, as shown in our test, does not constitute a good initial state. In fact, the overlap between our ground state ψ_0 and our *complete* initial state ψ_{init} was smaller than numerical accuracy even for only $L_b = 25$ bath sites. We conclude that the Fermi state constitutes a remarkably good guess, even at a relatively large number of bath sites. Note that, while in theory, the Fermi state is the exact solution solely for $U = 0$, we find that even for unreasonable large interaction strengths, the overlap is persistently high. The deviation from 1 for overlaps between our Fermi state and the ground state at $U = 0$ for large bath sizes originates from our construction principle as we restricted the bond dimension of our initial state to $m = 32$.

Both of these methods are an absolute necessity when performing 5 orbital simulations due to the exponential growth in complexity. Moreover, we want to mention that it is plausible that finite temperature simulations can become more efficient compared to zero-temperature simulations due to the infeasibility of finding the correct ground state quantum number sector. Especially since the energetic difference between quantum number sectors in multi-orbital systems tends to be smaller than errors encountered during time evolution. Hence, it is possible that fewer resources are needed to construct a good finite temperature state than to find the true ground state of the system.

5.4 Time evolution

We want to finish our discussion of the real time impurity solver with its, generally, most runtime-consuming part. Our quantity of interest is the impurity Green's function $G(\omega)$, which we can obtain via Fourier transformation from the retarded impurity Green's function:

$$G(\omega) = \int_{-\infty}^{\infty} e^{i(\omega+i\eta)t} G(t) dt, \quad (5.30)$$

where we added an additional Lorentzian broadening term η . The retarded impurity Green's function is given as

$$G(t) = -i\theta(t) \langle \psi_0 | \{ \hat{c}_i(t), \hat{c}_i^\dagger \} | \psi_0 \rangle, \quad (5.31)$$

where $|\psi_0\rangle$ denotes either the ground state of our impurity model for $T = 0$ simulations or a finite temperature state. We, again, shift the time evolution of the operator $\hat{c}_i(t)$ by the ground state energy E_0 as

$$\hat{c}_i(t) = e^{i(\hat{H}-E_0)t} \hat{c}_i e^{-i(\hat{H}-E_0)t}, \quad (5.32)$$

to compensate for the phase factor. We can simplify above equation by introducing the greater $G^>$ and lesser $G^<$ Green's functions

$$G(t) = -i\theta(t)(G^>(t) + G^<(t)), \quad (5.33)$$

where $G^>$ and $G^<$ are defined as

$$G^>(t) = \langle \Psi_0 | \hat{c}_i e^{-i(\hat{H}-E_0)t} \hat{c}_i^\dagger | \Psi_0 \rangle \quad (5.34)$$

$$G^<(t) = \langle \Psi_0 | \hat{c}_i^\dagger e^{i(\hat{H}-E_0)t} \hat{c}_i | \Psi_0 \rangle. \quad (5.35)$$

While a general discussion about tensor network methods would be no different to our presentation of the imaginary time solver, see section 4.3, there is an additional peculiarity of real time evolution of quantum states that we have to address, the entanglement growth.

In general, the entanglement growth in a quantum state during time evolution, quantified by the bipartite entanglement entropy S , see section 5.2, can be upper bounded as

$$S(t) \leq S(0) + a^{-1}vt, \quad (5.36)$$

where v is a typical velocity scale of the system, such as the Fermi velocity, and a a constant of unit length [6, 101–103]. This significantly limits the accessible times in tensor network simulations as the bond dimension can be related to the entanglement entropy with $m \sim e^S$ [102, 103]. Thus, this leads to a potentially exponential growth in computational cost. Note, however, that this is typically not the case for impurity models as we only consider local excitations where entanglement growth is typically logarithmic in time [104, 106].

Nonetheless, we want to address a strategy that can often be exploited to effectively reduce the entanglement growth during real time simulation. The homogeneity in time, which is typically found in impurity models, can be exploited to transform time-dependent correlation functions as [197, 198]

$$\langle \psi_0 | \hat{c}^\dagger(0) \hat{c}(t) | \psi_0 \rangle = \langle \psi_0 | \hat{c}^\dagger(-t') \hat{c}(t'') | \psi_0 \rangle, \quad (5.37)$$

where we have split t into $t = t' + t''$. We can rewrite the above expression as

$$\langle \psi_0 | \hat{c}^\dagger(0) \hat{c}(t) | \psi_0 \rangle = [\hat{c}(-t') | \psi_0 \rangle]^\dagger \hat{c}(t'') | \psi_0 \rangle. \quad (5.38)$$

Thus, the time evolution until t has been split into two time evolutions up to $-t'$ and t'' , respectively. Usually, entanglement growth is similar for *forward* and *backward* time evolutions. Hence optimal speed-ups can be achieved by symmetrically splitting t as $t' = t'' = t/2$. Thus, we trade a linear growth in the number of calculations we have to perform for a significantly reduced entanglement growth, see figure 5.21. Ultimately, this leads to a polynomial speed-up during time evolution. Note that this also significantly reduces errors during TDVP as it drastically lowers the truncation error due to the lower entanglement and further reduces the Trotter error due to the smaller time steps. However, reducing the number of time steps and keeping the time step fixed might be numerically favorable, as this further reduces projection and truncation errors during TDVP [92].




Figure 5.21: Effectiveness of splitting a regular real time evolution into a *forward* and *backward* evolution according to equation (5.38). We find a significantly reduced entanglement growth in (left), where we added an effective curve for the split method in which we double the time step to better visualize the stunted entanglement growth. We found the simulation using the split to be roughly 2-3 times faster compared to regular time evolution. In (right), we show the deviation of the retarded Green's functions for both methods. Data is shown for the SIAM model at half-filling, $U = 2D$ and $L_b = 59$.

Chapter 6

Multi-Orbital Impurity Solvers

In 2017, Bauernfeind et al. [39] published a new tensor network structure specifically optimized for multi-orbital impurity problems. They showcased its effectiveness based upon the simulation of real materials with up to 5 orbitals, a novelty at this point for tensor network based solvers. This work promised a paradigm shift in the feasibility of real time multi-orbital impurity solvers. However, it quickly became apparent that, though this network structure significantly outperforms MPS based solvers, it still suffers immensely from one of its core construction principles.

Baths of different orbitals typically do not interact directly with each other. Hence, a tensor network structure that ideally reflects the connectivity of the Hamiltonian would feature a connected impurity cluster with attached baths that only connect to their respective impurity site. Bauernfeind et al. proposed the Fork Tensor Product State (FTPS) network that fulfills all of these criteria; see figure 6.1. However, these specialized networks feature computationally unfavorable rank 4 tensors as impurity sites. Thus, updates involving two impurity sites scale as $\sim \mathcal{O}(m^5)$, marking them practically unfeasible for complicated real material simulations. The scaling of tensor network algorithms is roughly given as $\sim \mathcal{O}(m^{z+1})$ where z denotes the number of legs that need to be contracted for updates (marked as orange in figure 6.1).

Recently Gunst et al. [69] presented a novel idea of restructuring tree tensor network states (TTNS) by introducing branching nodes, nodes that solely feature virtual legs. This allowed them to maintain the favorable scaling of rank 3 tensors while preserving the representational power of tree tensor networks. We, thus, want to present novel tensor network structures, based upon these three-legged tree tensor network states (T3NS), to drastically improve the performance of tensor networked based multi-orbital impurity solvers. The general T3N framework has been created in close collaboration with Sam Mardazad; see [68].

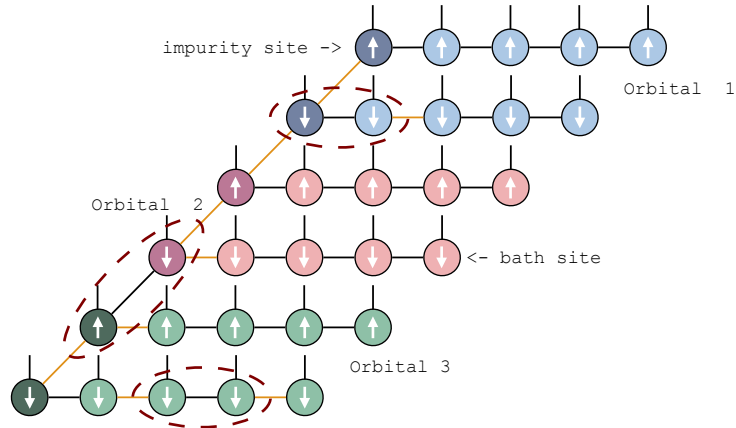


Figure 6.1: General structure of a FTSPS as it would be constructed for a three orbital model with separated spin degrees of freedom. Orbitals are marked with the same color, where impurity sites are shown in a darker shade, and bath sites are shown in a lighter tone of their respective color. Spin degrees of freedom are marked with an arrow. We find 3 different scalings for two-site updates on the FTSPS, dependent on the number of impurity sites involved in the update. For two impurity sites we find a scaling of $\sim \mathcal{O}(m^5)$ and for one and none a scaling of $\sim \mathcal{O}(m^4)$ and $\sim \mathcal{O}(m^3)$ respectively. Note that, due to the non-interacting nature of the bath sites, the impurity sites are typically significantly more entangled, resulting in the cost of variational updates of the network (DMRG or TDVP) being dominated by impurity sites.

6.1 Three-legged Tensor Network States

We will first establish multiple novel tensor network structures that are based on T3N but tailored towards multi-orbital impurity models. Afterward, we will provide benchmark results for them against our existing MPS solver and the current state-of-the-art FTSPS solver.

First, let us consider the full T3NS as proposed by Gunst et al. in 2018. Our version, which we adapted to impurity models, is constructed as a spanning tree of nodes that connects all impurity sites, which we will refer to as *base tree*. Each bath is represented by a tree of nodes connected to its respective impurity site via a branching node. This structure exponentially compresses the distance between bath sites. For a given number of layers N_L of our bath representation, we would have

$$L_b = \sum_k^{N_L} (Z - 1)^k \quad (6.1)$$

where Z is the coordination number of our network, i.e., $Z = 3$. Thus, while L_b grows exponentially with N_L , the distance between these sites still only depends linearly on N_L .




Figure 6.2: Depiction of a Three-Legged Tree Tensor Network (T3N) adapted for impurity models. Branching nodes are shown in grey, whereas impurity and bath sites are color-coded for each orbital, whereas impurity sites are shown in darker and bath sites in a lighter tone of their respective color. The *base tree* connecting all impurity sites is indicated by enlarged virtual bonds.

However, as we have found earlier, the correlation between bath sites is rather low in ideal tensor network representations, see section 5.2. Hence, this does not result in a significant reduction of entanglement in our system. In fact, as we shall find, the added number of branching nodes drastically increases the computational cost compared to simpler representations of the bath. This can simply be explained as single-site updates of branching nodes scale as $\sim \mathcal{O}(m^4)$ whereas updates of physical sites scale as $\sim \mathcal{O}(m^3)$. Thus, branching nodes must significantly lower the entanglement of the state. Otherwise, they will lead to additional costs in tensor network algorithms.

Hence, a generally better-suited tensor network structure is what we shall refer to as fork three-legged tree tensor network (FT3N). This name was deliberately chosen as it most closely compares to the FTPS. This structure allows for more efficient updates in the bath while keeping different baths spatially separated. This already constitutes a vast improvement over the full T3N due to the significantly lower number of branching nodes. Interestingly, ground state energies obtained using FT3N are only slightly worse than those obtained with T3N at a fixed bond dimension, see table 6.1. This indicates yet again that bath sites are only weakly coupled with each other. Hence, further reducing their distance from each other comes with little to no computational benefit.




Figure 6.3: Depiction of a fork three-legged tree tensor network (FT3N). Structurally, this bears the closest resemblance to the FT3N. However, note the additional branching node compared to the MT3N shown in figure 6.4.

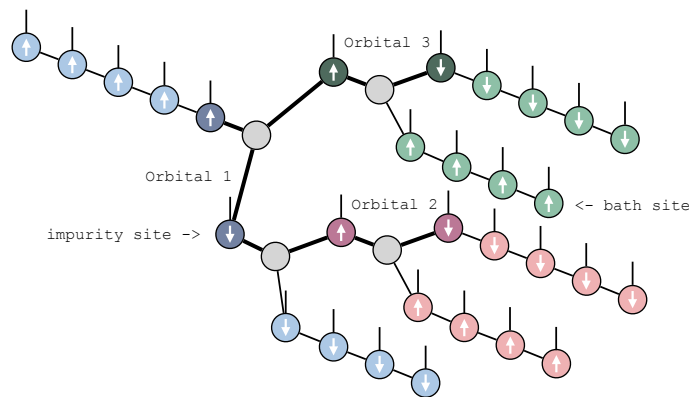


Figure 6.4: Depiction of a minimal three-legged tree tensor network (MT3N). The average distance between impurity sites is reduced compared to the FT3N by flipping one orbital and attaching it to the top of the tree. This figure has already been shown in [6]

The last structure we want to consider is the minimal three-legged tree tensor network (MT3N). The basic construction principle is simple: we want to create a network with a minimal number of branching nodes while maintaining the accustomed representation for multi-orbital impurity models. As we already stripped all branching nodes from our bath representations in the FT3N, a further reduction is only possible in the *base tree*. We can achieve this by inverting the order of one orbital and placing it on top of our tree, see figure 6.4. This saves one additional branching node. While this might appear negligible, updates involving branching nodes make up the lion’s share of computational cost during variational update procedures. Moreover, it reduces the average distance between impurity sites, which leads to a drastically improved convergence behavior.

We want to present benchmarks on the basis of a three-orbital model with a semi-elliptical density-of-states at half-filling and Hubbard-Kanamori interaction with $U = 2D$ and $J = 0.3D$ discretized with $L_b = 99$ bath sites per impurity. We disregard the $SU(2)$ symmetry in our simulations and use the full electronic basis. We find that all of our tree tensor networks significantly outperform our previous MPS impurity solver, some by as much as one order of magnitude in runtime. Note that we base the comparison on energy convergence per runtime as the direct comparison between bond dimensions in different tensor network structures is poorly defined. However, we kept the configuration of the local Krylov solver and a truncated weight of $w_t = 10^{-14}$ fixed for ground state searches. Note that we perform twice as many sweeps in our MPS simulations to reflect the typical tree tensor network sweeping pattern which would otherwise update sites twice as often [199].

Within our tree tensor network structures, we find that both the FT3N and MT3N show roughly the same performance, whereas the full T3N is about four times slower in comparison. However, we want to stress that the MT3N shows a significantly faster rate of convergence compared to the FT3N. In fact, the MT3N obtained the final, converged energy after just 1.5 hours, a result that the FT3N was unable to fully reproduce even after 4.26 hours. Therefore, we conclude our comparison by affirming that the MT3N represents the most effective T3N-based impurity solver. Moving forward, we will focus exclusively on comparisons involving the MT3N in the subsequent section. All benchmark simulations presented in this chapter have been performed on a single node with 2 Intel(R) Xeon(R) Platinum 8362 processors with 32 CPUs each.

Network	bond dimension	energy [eV]	runtime [h]
MPS	1024	-312.10878550	0.63
MPS	2048	-312.10885153	2.15
MPS	4096	-312.10886191	7.06
T3N	256	-312.10885192	0.96
T3N	512	-312.10886241	3.23
T3N	1024	-312.10886331	18.20
FT3N	256	-312.10885015	0.23
FT3N	512	-312.10886232	0.93
FT3N	1024	-312.10886330	4.26
MT3N	256	-312.10885768	0.21
MT3N	512	-312.10886305	1.15
MT3N	1024	-312.10886336	4.25

Table 6.1: Comparison of DMRG runtimes for different tensor network based impurity solvers. All considered tree-tensor networks clearly outperform the MPS based solver. Moreover, we want to stress that the MT3N solver reached the final energy after just 1.5 hours, thus also clearly outperforming the FT3N solver.

6.2 Comparison with state-of-the-art Tensor networks

We established that the MT3N is the most efficient tensor network representation for multi-orbital impurity models among our tested tree-tensor network structures. Now, let us provide comparisons based on full DMFT iterations with the current state-of-the-art tensor network structures, our MPS solver [2, 4, 6] and the FT3N solver [39, 200]. The FT3N simulations have been performed by Alexander Hampel [201].

Let us first establish the need for specialized tensor networks structures by presenting a comparison, yet again, on the basis of a simple multi-orbital impurity model at half-filling with a semi-elliptical density-of-states and Hubbard-Kanamori interactions, now with $U = D$ and $J = D/8$ and a split electronic basis. As a precise comparison of runtimes and associated errors is unfeasible, we present our best effort, which is to reach a similar precision among all our calculations. See table 6.2 for an overview of simulation parameters and table 6.3 for the associated mean sum rule errors $\delta = |n_{\text{GS}} - n_{\text{GF}}|$, which mark the difference between the mean ground state impurity occupations n_{GS} and the mean occupation derived from the impurity Green's functions n_{GF} . The latter can be obtained as

$$n_{\text{GF}} = - \int_{-\infty}^0 \text{Im} \frac{G(\omega)}{\pi} d\omega. \quad (6.2)$$




Figure 6.5: Runtime comparison of different tensor network structures for the imaginary time solver (left) and the real time solver (right). We find that, while all tensor network structures show an exponential growth in runtime with the number of orbitals, the base is significantly smaller for the MT3N, allowing for the description of full d- or even f-orbital system.

We resort to the mean as the error in multi-orbital impurity models in MPS greatly varies between impurities due to different spatial separations between impurity and bath sites.

All tensor network based impurity solvers show an exponential scaling in the number of impurity orbitals. However, the base is dramatically different leading to large overall speed-ups for the MT3N solver over our traditional MPS implementation. We find that the use of non-abelian symmetries becomes impactful only for our most difficult simulation of seven-orbital systems, see figure 6.5. However, we want to stress that real material simulations, especially finite temperature calculations, profit from the additional symmetry even at three-orbital simulations. Overall, we showcase the drastic difference in runtime requirements for real and imaginary time solvers. While five-orbital simulations with MPS are still feasible for the imaginary time solver, they are out of reach for the real time solver due to the larger exponential growth in runtime. Note that while it appears as if the five-orbital result for the real time solver would be even more accurate than the MT3N, this simply constitutes a failure of our measure as the low energy description of our model completely breaks down for the MPS and the small deviation from the sum rule is a mere coincidence. Moreover, with the advent of our MT3N solver, we showed that the simulation of seven-orbital impurity problems, i.e., f-orbital transition elements is in reach for us already, within constraints on energy resolution and orbital degeneracy, marking an exciting new research direction into their rich physics.

Network	solver	symmetry	$m_{N_{\text{orb}}=1}$	$m_{N_{\text{orb}}=2}$	$m_{N_{\text{orb}}=3}$	$m_{N_{\text{orb}}=5}$	$m_{N_{\text{orb}}=7}$
MPS	IT	U(1)	512	768	1024	1024	1536
MPS	IT	SU(2)	512	768	1024	1024	1200
MT3N	IT	U(1)	-	256	256	256	300
MPS	RT	U(1)	512	512	512	1024	-
MT3N	RT	U(1)	-	192	256	300	420

Table 6.2: Bond dimensions during time evolution used for our scaling comparison of different tensor network solvers, see figure 6.5. All imaginary time simulations have been performed at $\beta_{\text{eff}} = 400$ with $L_{\text{b}} = 9$ bath sites and a truncated weight of $w_t = 10^{-11}$. We used $L_{\text{b}} = 199$ bath sites and $\eta = 0.005D$ for our real time solver. Again, time evolution was performed with a truncated weight of $w_t = 10^{-11}$.

Network	solver	symmetry	$\delta_{N_{\text{orb}}=1}$	$\delta_{N_{\text{orb}}=2}$	$\delta_{N_{\text{orb}}=3}$	$\delta_{N_{\text{orb}}=5}$	$\delta_{N_{\text{orb}}=7}$
MPS	IT	U(1)	0.0004	0.0007	0.0007	0.0013	0.0157
MPS	IT	SU(2)	0.0004	0.0007	0.0007	0.0008	0.0013
MT3N	IT	U(1)	-	0.0006	0.0007	0.0009	0.0010
MPS	RT	U(1)	0.0020	0.0025	0.0046	0.0024	-
MT3N	RT	U(1)	-	0.0025	0.0029	0.0034	0.0040

Table 6.3: Mean sum rule error $\delta = |n_{\text{GS}} - n_{\text{GF}}|$. This measure breaks down in the presence of stark oscillations as they do not affect the total weight. While the 5-band MPS simulation of our RT solver appears more accurate than the MT3N solver, the contrary is the case. The accuracy of the MPS solver at small frequencies breaks down completely and the better conservation of the sum rule is just a mere coincidence.

So far, the FTPS constituted the most successful tensor network impurity solver. However, as established earlier, it suffers from a fundamental flaw, the presence of rank 4 tensors, see figure 6.1. We base our comparison on the simulation of the t_{2g} orbitals of the insulating perovskite SrMoO_3 [200]. We find an overall good agreement between both solvers; see figure 6.6. Note that the slight deviation between our results can be explained by the unphysical behavior of the FTPS solver at high energies. The real part of the self-energy should become stationary for higher energies. However, likely due to a faulty Fourier transformation, this behavior is not observed in the FTPS result.




Figure 6.6: Comparison between self-energies obtained with FTPS and our novel MT3N structure. We find an overall good agreement between both solvers. However, note that the deviation at higher energies is likely the result of a faulty Fourier transformation as the real part of the FTPS self-energy does not approach the stationary limit.

We present runtime comparisons between the FTPS solver and MT3N in table 6.4. Time evolution was performed until $t_{\max} = 70.6\text{eV}^{-1}$ with a maximum bond dimension of $m = 200$ and a truncated weight of $w_t = 10^{-10}$. Both solvers used $L_b = 229$ bath sites to discretize the hybridization function using the fit discretization method with $\eta = 0.08$. The FTPS solver was unable to complete time evolution within the maximally allowed time frame of 5 days using 2TDVP due to its unfavorable $\mathcal{O}(m^5)$ scaling. The MT3N solver is nearly as fast as the FTPS 1TDVP runtime using 2TDVP and posts an impressive ~ 2.8 times speed-up over the FTPS using 1TDVP. We present a further comparison with an additional TDVP mode which we dub 1.5TDVP. In this mode, we only perform single-site updates but contract branching nodes with the next physical node in sweeping direction prior to the update, effectively yielding rank 4 tensors, i.e. a local correspondence to FTPS.

The large speed-up of the MT3N over the FTPS in 1TDVP can be roughly explained by counting the number of $\mathcal{O}(m^4)$ updates. Considering the unfolded representation shown in figure 6.1, we find 4 rank 4 tensors, as the outer two impurity sites are only rank 3 tensors. At a first glance, the MT3N also has 4 $\mathcal{O}(m^4)$ updates. However, branching nodes do not have physical legs. Thus, we would expect the MT3N to be faster by roughly a factor of d where d is the local physical basis. Indeed this is roughly the speed up we find in our simulations. In theory, the 1.5TDVP update scheme should match the performance of the FTPS as branching nodes are pre-contracted with physical nodes for updates, thus resulting in updates on rank 4 tensors. We hence conclude that we have an additional ~ 1.5 times speed-up that is simply the result of individual implementation details between the ITensor [202, 203] based FTPS solver and our SYTEN[77] toolkit.

Network	1TDVP [h]	1.5TDVP [h]	2TDVP [h]
FTPS	6.6	-	-
MT3N	2.3	3.65	8.65

Table 6.4: Runtime comparison for different TDVP modes for the FTPS and MT3N structures.

Chapter 7

Complex Time Impurity Solver

The following results are based on the author's publication [6].

Real time impurity solvers are plagued by a growing complexity in two quantities. First, the exponential growth of computational cost in the number of orbitals which we addressed with our novel MT3N structure that significantly lowers the base of said growth, see figure 6.5. However, the arguably more fundamental challenge of the RT solver is the polynomial growth in complexity during real time evolution, see section 5.4. This strongly restricts the achievable frequency resolution of Green's functions and, subsequently, self-energies and other derived quantities as the long time limit is generally inaccessible to tensor network methods. By extending the time evolution to contours in the complex plane, entanglement growth is curtailed, enabling numerically efficient high-precision calculations of time-dependent correlators and Green's functions with detailed frequency resolution.

Our goal is to combine real and imaginary time evolutions in the complex plane to develop methods to extract real-time information while using imaginary-time evolution to limit entanglement growth during time evolution. Imaginary time-evolution methods have been frequently used for ground-state searches, which is reflected by the fact that imaginary time evolution generally acts as a gradual energy truncation: it suppresses high-energy states and enhances the weight of low-energy states exponentially in the long-time limit [90, 204]. So far, energy truncation algorithms have only been combined successfully with tensor network methods that calculated correlation functions directly in frequency space [98, 205, 206]. This leads to much faster, less resource-intensive calculations, which in turn give cheaper access to the same information as real-time evolutions or allow one to proceed to longer times and with better frequency resolution. The idea of complex time evolution has been recently discussed in the context of Quantum Monte Carlo simulations [207]. Similar ideas are also published in another work [208].

We introduce contours in the complex plane as $z(t) = t + i\tau(t)$ and define a general

time-dependent correlation function $C_{OP}(z)$ as

$$C_{OP}(z) = \langle \psi_0 | \hat{O}^\dagger(z) \hat{P} | \psi_0 \rangle, \quad (7.1)$$

where $|\psi_0\rangle$ is the ground state with energy E_0 and

$$\hat{O}(z) = e^{+i(\hat{H}-E_0)z} \hat{O} e^{-i(\hat{H}-E_0)\bar{z}}, \quad (7.2)$$

with \bar{z} the complex-conjugate of z . Effectively, then

$$C_{OP}(z) = \langle \psi_0 | \hat{O}^\dagger e^{-i(\hat{H}-E_0)\bar{z}} \hat{P} | \psi_0 \rangle. \quad (7.3)$$

Let us quickly visualize the core principle of complex time evolution before we move on to an in-depth discussion of different contours in the complex-time domain. For the sake of simplicity, let us consider a free electron placed at the center of a lattice with 200 sites. The electron is governed by the usual tight binding model dispersion relation $E(k) = -2t \cos(ka)$ with $t = a = 1$. We time evolve it along a complex contour tilted by various angles α with respect to the real axis, where $z(t) = t + it \tan \alpha$, i.e. $\tau = t \tan \alpha > 0$, see section 7.2 for an in-depth discussion of this contour and its properties. We observe that, upon time evolution in imaginary time direction, the energy decreases towards its minimum of -2 as high energy contributions are damped out as we continuously apply powers of our Hamiltonian, akin to an energy truncation. This contrasts the purely unitary real time evolution that conserves the energy, see figure 7.1. This has two major beneficial effects. First, entanglement is drastically reduced as ground states of gapped Hamiltonians typically follow an *area law* entanglement [74], contrary to the *volume law* potentially encountered for bulk states [209]. Secondly, it limits the extent of the light cone, which in turn, again, potentially reduces the entanglement growth based on the upper bound that can be placed upon the growth of the bipartite entanglement entropy S during real time evolution of a local excitation

$$S(t) \leq S(0) + a^{-1}v \log t, \quad (7.4)$$

where v is the phase velocity in our simple example [101–103]. Note that we have only considered the light cone spread in a single particle picture. A full investigation of many particle systems with different dispersion relations and interactions is still needed to conclusively address the effect of complex time evolution on the spread of information in quantum many-body systems. However, we want to stress that this is a remarkable property as it can make time evolutions of complicated two-dimensional or open quantum systems accessible to tensor network methods.




Figure 7.1: Decrease of the normalized energy expectation value over time on tilted complex contours with angles α in a). We calculate the phase velocity from the dispersion relation in b) at energies realized at different times, marked as black dots in a). We plot the phase velocity at the same points in time, again marked with a black dot, alongside the light cone spread in c) indicated by contour plots of the block entropy. We observe that the extracted phase velocities from b) align perfectly with the observed light cone in c) validating our simple picture for complex time evolution.

This chapter is structured as follows: first, we present easily implemented contours for which reliable post-processing methods exist to extract the desired real-time result. The various schemes allow mutual verification and present a different balance of speed vs precision. Then we will present comparisons both for the impurity spectral function and for the numerically challenging calculation of the self-energy at very low frequencies, down to the Fermi liquid regime where $\text{Im} \Sigma(\omega) \sim \omega^2$. We base these comparisons on the single-impurity Anderson model for which highly precise data is available, both from tensor network methods and the numerical renormalization group (NRG). Finally, we will illustrate the potential of our method on the example of a model with three impurity orbitals that interact via Hubbard-Kanamori or Dworin-Narath interactions.




Figure 7.2: Parallel contour $t + i\tau$. Multiple complex parallel contours are obtained with a small shift $\delta\tau$ from a first contour, cf. text.

7.1 Constant Imaginary Time: Parallel Contour

We now turn to the first complex-time method, where time is evolved along a complex contour parallel to the real-time axis, i.e., at constant imaginary time, see figure 7.2. We have $z(t) = t + i\tau$ with $t \in [0, t_{\max}]$ and $\tau > 0$ constant. Here, z starts at 0, moves up to $i\tau$, and then continues from $i\tau$ to $t_{\max} + i\tau$. The entanglement growth with time is strongly limited compared to the real time evolution, as illustrated in figure 7.3. This contour suppresses high energy states aggressively, as reflected in the small entanglement growth at short times compared to other contours, see figure A.1, where we compare the time dependency of the entanglement for all the contours introduced in this work. It is, therefore, advantageous to use comparatively small $\tau > 0$ to retain a sufficient amount of high energy spectral weight when used in combination with post-processing methods. An attractive feature of the parallel contour is that additional parallel contours offset by $\delta\tau$ can be obtained very cheaply from a first contour. Let us illustrate this point with one correlator

$$\langle \psi_0 | \hat{c}^\dagger \hat{c}(t + i(\tau + \delta\tau)) | \psi_0 \rangle = [e^{-(\hat{H} - E_0)\delta\tau} \hat{c} | \psi_0 \rangle]^\dagger e^{i(\hat{H} - E_0)(t + i\tau)} \hat{c} | \psi_0 \rangle. \quad (7.5)$$

We first compute the time evolution of the right ket up to time $t + i\tau$, and the additional time evolution $\delta\tau$ for the left bra, which is cheap to compute for small $\delta\tau$. Note that we can use positive or negative $\delta\tau$. In practice, we use $\delta\tau < 0$, as illustrated in figure 7.2. This allows us to obtain the correlations on a contour closer to the real axis, which is a priori harder to compute, at the price of a positive exponential evolution of the left term for a short time $\delta\tau$.

Let us introduce the complex-time spectral function \mathcal{A}_τ defined by

$$\mathcal{A}_\tau(\omega) \equiv \frac{1}{2\pi} \int dt e^{i(\omega + i\eta)t} \langle \psi_0 | \{ \hat{c}_0(t + i\tau), \hat{c}_0^\dagger \} | \psi_0 \rangle. \quad (7.6)$$




Figure 7.3: Evolution of the bond dimension profiles during complex time evolution (parallel contour at $\tau = 1.3$) in (left) and for real time evolution in (right). Both evolutions were obtained for the SIAM model at $U/D = 2$ with $L_b = 199$ bath sites with a basis that has been split into spin-up (sites 0-199) and spin-down (sites 200-399) degrees of freedom. We find a significantly slower growth of bond dimension for the complex time evolution, note the different time slices. It becomes obvious that imaginary time evolution acts as an energy truncation as sites are ordered in energy. As this is a particle-hole symmetric system, sites around 100 and 300 have on-site potentials that are nearly zero. Thus these are the only sites that show significant correlation. We find that the real time evolution in (right) quickly saturates the maximum bond dimension of $m = 1024$.

At $\tau = 0$, this function is the ordinary spectral function $\mathcal{A}_{\tau=0}(\omega) = \mathcal{A}(\omega) = -\frac{1}{\pi} \text{Im} G(\omega)$. Due to the simple form of the contour, we can derive an explicit relation between \mathcal{A}_τ and \mathcal{A} :

$$\mathcal{A}_\tau(\omega) = \mathcal{A}(\omega)e^{-\tau|\omega|}, \quad (7.7)$$

which is exact in the limit $\eta \rightarrow 0$ as is used here. Using the definition of the complex time evolution equation (7.2), we obtain

$$\begin{aligned} \langle \psi_0 | \{ \hat{c}_0(z(t)), \hat{c}_0^\dagger \} | \psi_0 \rangle &= \sum_A |\langle \psi_0 | \hat{c}_0 | A \rangle|^2 e^{-iE_A(t-i\tau)} \\ &\quad + |\langle \psi_0 | \hat{c}_0^\dagger | A \rangle|^2 e^{+iE_A(t+i\tau)}, \end{aligned}$$

where $|A\rangle$ is a eigenstate basis of the many-body Hamiltonian $H - E_0$, and E_A its eigenvalue. From the definition of \mathcal{A}_τ (7.6), we have

$$\begin{aligned} \mathcal{A}_\tau(\omega) &= \sum_A |\langle \psi_0 | \hat{c}_0 | A \rangle|^2 \delta(\omega - E_A) e^{-E_A \tau} + \\ &\quad |\langle \psi_0 | \hat{c}_0^\dagger | A \rangle|^2 \delta(\omega + E_A) e^{-E_A \tau} \end{aligned}$$

in the limit $\eta \rightarrow 0$. By definition of the ground state, $E_A > 0$. In the first term, $\omega = E_A = |\omega|$, while in the second term $E_A = -\omega = |\omega|$, so we get

$$\mathcal{A}_\tau(\omega) = e^{-|\omega|\tau} \sum_A |\langle \psi_0 | \hat{c}_0 | A \rangle|^2 \delta(\omega - E_A) + |\langle \psi_0 | \hat{c}_0^\dagger | A \rangle|^2 \delta(\omega + E_A)$$

and therefore

$$\mathcal{A}_\tau(\omega) = \mathcal{A}_{\tau=0}(\omega) e^{-|\omega|\tau}.$$

Our method consists simply in computing $\mathcal{A}_\tau(t)$, then $\mathcal{A}_\tau(\omega)$ using a Fourier transform, and finally invert (7.7) to obtain the spectral function $\mathcal{A}(\omega)$. Despite its apparent simplicity, this inversion presents, however, two difficulties, both at low and high frequencies.

First, the inversion of (7.7) is clearly difficult at high frequencies, as small errors in $\mathcal{A}_\tau(\omega)$ are amplified by the exponential. Such errors may result from a too small t_{\max} for $\mathcal{A}_\tau(t)$ to be fully decayed, or from a broadening in the Fourier transform. This issue can be solved by introducing a cut-off frequency ω_c that limits the growth of the exponential factor as

$$\mathcal{A}(\omega) = \mathcal{A}_\tau(\omega) e^{\tau \min(|\omega|, \omega_c)}, \quad (7.8)$$

where ω_c is chosen such that it only acts in the high-frequency tail of the spectral function. We found $\omega_c = 3D$ to yield overall good results in our calculations.

Second, the inversion of (7.7) is also difficult at low frequencies, even though the exponential term is close to 1. In DMFT, and in many other physics applications, we are interested in a high-precision computation of the behavior of the spectral function $\mathcal{A}(\omega)$ or the self-energy $\Sigma(\omega)$ at low frequencies. The difficulty comes from the kink in $\mathcal{A}_\tau(\omega)$ at $\omega = 0$ in Eq. (7.7), caused by $|\omega|$ in $e^{-\tau|\omega|}$. This is an exact feature, but as $\mathcal{A}(\omega)$ does not exhibit a kink at $\omega = 0$, this must be exactly compensated by a kink in $\mathcal{A}_\tau(\omega)$ at $\omega = 0$. The latter requires high-accuracy results from long-time evolutions where tensor network methods are limited. As a result, the compensation is imperfect; see figure 7.4.

Our solution to this issue uses a linear combination of a few $\mathcal{A}_{\tau_k}(\omega)$ computed on parallel contours $t + i\tau_k$ (as described above), which is designed to compensate for the low-frequency singularity introduced by the $e^{-\tau|\omega|}$ factor. Introducing some weights a_k , and defining the function

$$h(\omega) \equiv \sum_{k=1}^n a_k e^{-\tau_k |\omega|}, \quad (7.9)$$

we have

$$\mathcal{A}(\omega) h(\omega) = \sum_{k=1}^n a_k \mathcal{A}_{\tau_k}(\omega). \quad (7.10)$$

We choose the weights a_k such that the function $h(\omega)$ is flat close to $\omega = 0$, see figure 7.4,




Figure 7.4: Comparison of $h(\omega)$ for n equally spaced τ_k between $[1.3 - 2(n-1)\delta\tau, 1.3]$ with $\delta\tau = 0.075$ (left) and the corresponding spectral functions (right) of the SIAM model with $U/D = 2$, $N_b = 299$ and $t_{\max} = 150D^{-1}$. One can clearly see the effect of the kink in the vicinity of $\omega = 0$ for the direct application of equation (7.7). Our linear superposition approach practically cancels the effect of the kink while maintaining an overall high accuracy and a good agreement within different orders of the scheme with the exception of $n = 5$.

i.e. we cancel the first powers of its low ω expansion:

$$\sum_{k=1}^n a_k = 1 \quad (7.11a)$$

$$\sum_{k=1}^n a_k \tau_k^l = 0 \quad (7.11b)$$

for $l = 1, \dots, n-1$. As h is close to 1 at low ω , $h(\omega) = 1 + O(\omega^n)$, we can now safely invert (7.10):

$$\mathcal{A}(\omega) = \frac{1}{h(\omega)} \sum_{k=1}^n a_k \mathcal{A}_{\tau_k}(\omega). \quad (7.12)$$

In practice, we need to take τ_k , which are neither too close (leading to large a_i that amplify numerical noise, due to the Vandermonde determinant in the linear system (7.11)), nor too distant as time-evolution errors affect precision. For reasonably selected a_i , the overall result does not strongly depend upon n , see figure 7.4. In this work, we used $n = 3$ and $\tau_k = 1, 1.15, 1.3$, if not stated otherwise. We will refer to this approach as the parallel (inversion) method.

A direct extrapolation, which we call the parallel (extrapolation) method, to $\tau = 0$ using several parallel contours can also be performed. We present it in appendix A.1 for completeness, but it is slightly inferior to the parallel (inversion) method. The parallel (extrapolation) method will hence be used as a check for the parallel (inversion) method.

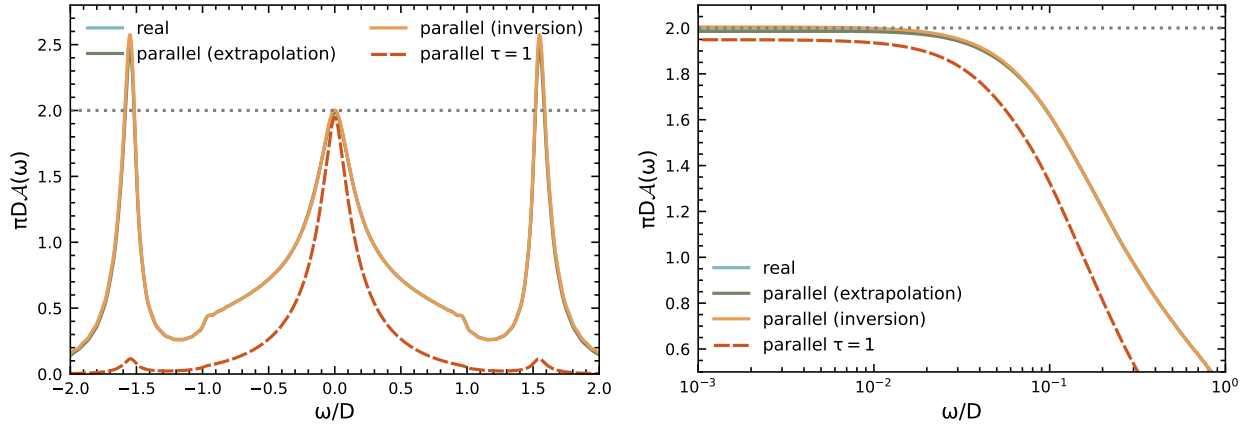


Figure 7.5: Comparison of spectral functions obtained on a parallel contour (left): "real" stands for the real-time reference data, $\tau = 1$ on the parallel contour without post-processing, parallel (extrapolation), and parallel (inversion) for the two methods mentioned in the main text. Nearly all spectral weight can be recovered by both methods in excellent agreement. The parallel (inversion) method used $n = 3$ with $\tau_k = 1, 1.15, 1.3$; the coefficients a_k (see text) are then approximately 33.2, -57.8 , 25.6. The parallel (extrapolation) method used corrections to 6th order with 13 contours centered on $\tau = 1$ with a distance $\delta\tau = 0.075$. Zoom at low frequencies with the same labels (right), showing the high accuracy of the parallel (inversion) and parallel (extrapolation) methods at very low frequencies.

We present results of the spectral function of the benchmark SIAM in figure 7.5. We use $\delta t = 0.2D^{-1}$, $w_t = 10^{-11}$, $N_b = 59$, $t_{\max} = 90D^{-1}$, and $\eta = 0.001D$. The simulation reached a maximum bond dimension of $m = 1023$. Note that the simulations using purely real time evolution resulted in a maximum bond dimension of $m \sim 1500$ for $w_t = 10^{-10}$ and $m \sim 2700$ for $w_t = 10^{-11}$. Both the parallel (inversion) and parallel (extrapolation) methods provide spectral functions in excellent agreement with the much more costly (more than an order of magnitude) real-time benchmark data both at higher and, in particular, also at very low frequencies. The (generalized) Friedel sum rule $\pi D\mathcal{A}(0) = 2$ is matched to fractions of a percent. As both methods operate on the same numerical data, the parallel (extrapolation) method provides a cheap control of the quality of the spectral function. For figure 7.6, we use larger baths and a larger maximum bond dimension to calculate the spectral function for larger values of U/D both by real-time calculations and by the parallel (inversion) method. The Friedel sum rule is matched to very high accuracy in all cases; real-time results are much less accurate for larger U/D .

The determination of the self-energy $\Sigma(\omega)$ is numerically more challenging than that of the spectral function $\mathcal{A}(\omega)$. We calculate it from the Dyson equation $\Sigma(\omega) = G_0^{-1}(\omega) - G^{-1}(\omega)$, where $G_0(\omega)$ is the impurity Green's function for the non-interacting case. Its inverse is given by $G_0(\omega)^{-1} = \omega + i\eta - \epsilon_0 - \Delta(\omega)$ with $\eta = 0$ here. We used the analytically known




Figure 7.6: Spectral function of the SIAM for $U/D \in \{2, 3, 4\}$ with $N_b = 299$, using the $n = 2$ parallel (inversion) method with $\tau_k \in \{1.15, 1.3\}$ (left); the coefficients are approximately $a_k \in \{8.6, -7.6\}$. We have an overall higher frequency resolution due to a finer bath discretization, zero broadening, and a time evolution $t_{\max} \in \{150D^{-1}, 200D^{-1}, 250D^{-1}\}$. The maximum bond dimension is $m = 1500$ $SU(2)$ states. In all cases, in particular for larger U/D , the matching of the Friedel sum rule is less accurate for the real-time method. The corresponding self-energies are shown in (right).

hybridization function. We obtain $G(\omega)$ from the parallel (inversion) method, which yields $\mathcal{A}(\omega)$, i.e. the imaginary part of $G(\omega)$. We then use the Kramers–Kronig rule to obtain the real part of $G(\omega)$.

To assess the quality of our self-energies, we henceforth include benchmark results from NRG. The NRG data was obtained in a state-of-the-art implementation [123, 210–214] based on the QSpace tensor library [215], using a symmetric improved estimator for the self-energy [179]. All NRG simulations have been performed by Fabian Kugler. The latter allows one to follow the imaginary part of the self-energy down to extremely low values of $|\text{Im } \Sigma|/D$.

As shown in figure 7.6, the self energies calculated by complex time evolution reach the Fermi-liquid ω^2 -regime. The final breakdown at very low frequencies ($\omega/D \approx 0.007$) is due to small deviations from the sum rule and the use of Dyson’s equation. We expect that the use of improved estimators [179] will alleviate this problem in future implementations. Before the breakdown, agreement with NRG results is excellent. Note that self-energies calculated from real-time results are much less accurate and plagued by unphysical oscillations.




Figure 7.7: Symmetrically split complex-time contour: Contour \mathcal{C}_{res} from 0 to z is replaced by two symmetric contours \mathcal{C}_+ and \mathcal{C}_- from 0 to $z/2$ and $-\bar{z}/2$ respectively.

7.2 Time Evolution at Fixed Angles: Tilted Contour

In our second complex-time approach, time is evolved along a complex contour tilted by various angles α with respect to the real axis, where $z(t) = t + it \tan \alpha$, i.e. $\tau = t \tan \alpha > 0$. We adapt the symmetric splitting of real-time evolutions to the complex plane by splitting $z = z' + z''$ with $z' = z'' = z/2$. Then we have

$$\langle \psi_0 | \hat{c}^\dagger(0) \hat{c}(z) | \psi_0 \rangle = [\hat{c}(-\bar{z}/2) | \psi_0 \rangle]^\dagger \hat{c}(z/2) | \psi_0 \rangle, \quad (7.13)$$

$$\langle \psi_0 | \hat{c}(z) \hat{c}^\dagger(0) | \psi_0 \rangle = [\hat{c}^\dagger(z/2) | \psi_0 \rangle]^\dagger \hat{c}^\dagger(-\bar{z}/2) | \psi_0 \rangle. \quad (7.14)$$

The two contours on which complex time evolution occurs now look as in Fig. 7.7. In all calculations, we use $\delta|z| = 0.2D^{-1}$, and again $w_t = 10^{-11}$ and $N_b = 59$. The maximum complex time is in all cases $|z|_{\text{max}} = 90D^{-1}$, or real times $t_{\text{max}} = 90D^{-1} \cos \alpha$. The resulting Green's function is Fourier transformed as

$$G_\alpha(\omega) = \int_0^\infty dt G_\alpha(t + i\tau) e^{i\omega t - \eta t}, \quad (7.15)$$

again with $\eta = 0.001D$. The subscript α indicates that we evaluate on a tilted contour. The spectral function $\mathcal{A}(\omega)$ is extracted as $\mathcal{A}_\alpha(\omega) \equiv -(1/\pi)\text{Im}G_\alpha(\omega)$. Results are shown on a linear and a logarithmic frequency scale in figure 7.8. We do not expect to obtain the correct $\mathcal{A}(\omega)$ for $\alpha \neq 0$, as \mathcal{A}_α is just an intermediate step in the computation. Nevertheless, we find that peak positions in \mathcal{A}_α are well preserved for moderate angles, thus still allowing for a rough interpretation of quasi-particle peaks. We report a significant reduction in the entanglement entropy for all angles that can be understood as the result of damping the statistical weight of high energy states, as reflected in a reduction in the time-dependent energy expectation value of the system, see figure 7.9. The dampened growth in the entanglement entropy leads to a speed-up of a factor of the order 100 in the SIAM model for the specified parameters compared to our real-time reference calculation. This is reflected in the very small final bond dimensions, $m = 52$ for $\alpha = 0.3$, $m = 66$ for $\alpha = 0.2$, and $m = 137$ for $\alpha = 0.1$.

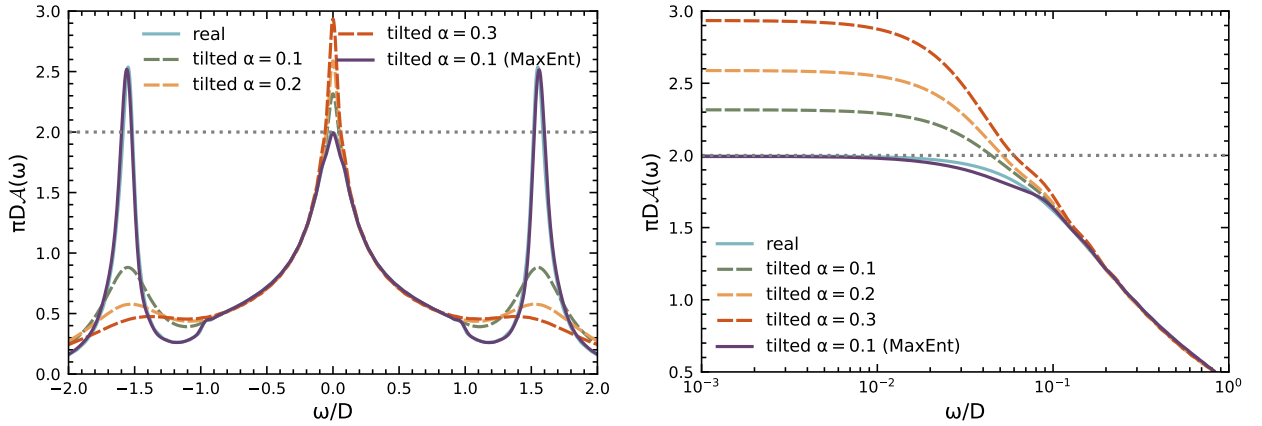


Figure 7.8: Comparison of spectral functions obtained on tilted contours with real-axis reference data ('real') and the tilted (MaxEnt) result from the tilted contour with $\alpha = 0.1$ (left). The kink at $\omega/D = \pm 1$ is model-specific. The dashed lines are data before post-processing. We show a zoom the spectral functions of (left) on a logarithmic frequency scale for $\omega/D > 0$ in $[10^{-3}, 10^0]$ in (right).




Figure 7.9: Energy (left) and entanglement entropy (right) during time-evolution for $U/D = 2$. Imaginary time evolution along a tilted contour shifts the energy expectation value of the time-evolving state towards the ground state energy. Entanglement entropy, here given by the sum of the entanglement entropies for all system cuts, is reduced strongly along the tilted contours, resulting in substantial speed-ups.

To compute the spectral function $\mathcal{A}(\omega)$, the data has to be post-processed, in this case by an analytical continuation of the complex-time data to the real axis by MaxEnt, yielding the tilted (MaxEnt) method. We follow the procedure outlined in [207] for the analytical continuation. There is a connection between $G(z)$ and the spectral function \mathcal{A} in the case of complex-time contours, given by

$$G(z) = -i\Theta(t) \int_{-\infty}^{\infty} d\omega \mathcal{A}(\omega) K(z, \omega), \quad (7.16)$$

where the integration kernel $K(z, \omega)$ is defined as

$$K(z, \omega) = \begin{cases} \exp(-i\bar{z}\omega), & \text{if } \omega \geq 0 \\ \exp(-iz\omega), & \text{if } \omega < 0. \end{cases}$$

As the kernel may be complex-valued, the equations have to be slightly adapted [207]. Besides this, the well-established methods of MaxEnt were used [152, 216], see also section 4.4. Analytic continuation was rather stable with this kernel. However, care must be taken at low frequencies as the kink of the kernel at zero frequency causes results to deviate from spectral functions obtained from purely real-time evolution. We propose two different measures as possible solutions. First, we use a broadening around the discontinuity by replacing the problematic $\text{sgn}(\omega)$ with $\tanh(\omega/\sigma)$, with the free factor σ . The kernel then reads

$$K(z, \omega) = \exp\left(- (it + \tanh(\omega/\sigma)\tau)\omega\right) \quad (7.17)$$

In the limit $\sigma \rightarrow 0$, this kernel is exact. Hence, σ has to be chosen sufficiently small to leave the overall structure of the inversion problem unchanged. We found sensible values to be smaller than the width of the main peak at $\sigma = 0.11$. Although this change helps a bit, it did not yield reliable results, leading to strong oscillations around $\omega = 0$. We therefore implemented a second method, which introduces a correction term to $Q(\mathcal{A}) = -\chi^2 + S(\mathcal{A})$, that is maximized in normal MaxEnt. We rather take

$$Q(\mathcal{A}) = -\chi^2 + \alpha S(\mathcal{A}) - \beta \int_{-\infty}^{\infty} d\omega \left(\frac{d\mathcal{A}}{d\omega}\right)^2 f(\omega). \quad (7.18)$$

The additional last term favors smooth spectral functions since, when maximizing Q , it tries to minimize the quadratic slope of \mathcal{A} . The observed quick oscillations around zero are thus smoothed out as they increase the quadratic slope. $f(\omega)$ is a weighting factor focusing on small frequencies. We used a Gaussian $f(\omega) = \exp(-(\omega/\sigma)^2)$. The new parameter β controls this correction term; we found $\beta \approx \alpha$ to be a viable choice. For all MaxEnt calculations involving the three-band model, we chose $\sigma = 1$. The MaxEnt with correction term maintains the correct peak heights. Finding an even better stabilization of the Kernel at low frequencies is subject of ongoing research.




Figure 7.10: Contour for complex time evolution with no post-processing (left). Stable version with a horizontal contour, cf Text (right).

The agreement with the real time data is overall very good, and the Friedel sum rule is obeyed to 0.4%, yet there are small but noticeable deviations from the true $\mathcal{A}(\omega)$ between $\omega/D \sim 10^{-2}$ and $\omega/D \sim 10^{-1}$ (figure 7.8). This is not surprising as the analytic continuation kernel is generally ill-conditioned, and its current formulation leads to deviations around the Fermi edge. In comparison to the parallel (inversion) and parallel (extrapolation) methods, accuracy is lower for the tilted (MaxEnt) method, in particular at low frequencies. The tilted (MaxEnt) method is, however, the fastest, so it may be very useful in DMFT applications for the intermediate steps of the iteration procedure. Note that the tilted contour can also be combined with the extrapolation scheme of appendix A.1 to yield the tilted (extrapolation) method; results will be shown for the self-energies calculated in the following section.

7.3 Real Frequency Results Without Analytic Continuation

For an additional approach, we reformulate the idea of time splitting in the complex plane in a different way by splitting *real* time into two *complex* times via $t = z' + z''$ with $z' = t' - i\tau$, $z'' = t'' + i\tau$, $t, t', t'', \tau > 0$ where $t = t' + t''$. Then

$$\langle \psi_0 | c^\dagger(0)c(t) | \psi_0 \rangle = [c(-z')] | \psi_0 \rangle]^\dagger c(z'') | \psi_0 \rangle. \quad (7.19)$$

In this case, no analytical continuation or any other post-processing is required.




Figure 7.11: Imaginary part of the retarded Green's function vs t and spectral function vs ω for the SIAM at $U/D = 2$ and other parameters as before, obtained using the kink contour method with $\alpha'' = 0.1$ and $\alpha' = \frac{\pi}{2}$. The "stable" label refers to the second contour in figure 7.10 with $t_{\max} = 20D^{-1}$ or $\tau_{\max} = 0.998 = 1.996D^{-1}$. The "stable" contour result is in excellent agreement with the real time reference data.

If we continue to use linear contours in the complex plane, we get two contours \mathcal{C}_+ and \mathcal{C}_- in the upper right and lower left quadrants of the complex plane at angles α' and α'' which need not be identical as no symmetric splitting of t is required (see figure 7.10). The two angles α', α'' can be chosen freely, but the imaginary time steps in \mathcal{C}_+ suppress the entanglement growth of the real-time evolution, whereas the same steps increase that entanglement growth on contour \mathcal{C}_- . This suggests a splitting where $\alpha' > \alpha''$ and hence $t' < t''$. In our calculations, we chose $\alpha'' = 0.1$ and $\alpha' = \pi/2$, i.e. a purely imaginary time evolution on a vertical contour \mathcal{C}_- ; $t = t''$ and $t' = 0$.

Numerical instabilities in the time evolution on \mathcal{C}_- limit our maximum evolution time on the corresponding contour \mathcal{C}_+ to about $|z''| = 20D^{-1}$ or $\tau'' = 1.996D^{-1}$ for $\alpha'' = 0.1$ and $\alpha' = \pi/2$. The numerical instabilities in the time evolution on \mathcal{C}_- can be seen clearly in figure 7.11. It might be suspected that the exponential growth of the weight of higher-energy states on contour \mathcal{C}_- suppresses the contribution from low-energy states, which are enhanced on contour \mathcal{C}_+ and numerical cancellations fail, leading to the instabilities. At the onset of the instabilities, the norm difference is only of the order 10^4 , which should only lose insignificant digits. So even though exponential growth will eventually make this method unstable, it does not seem to be the origin of the currently observed instabilities. We suspect that they are related to the relatively subtle interplay of errors in the TDVP method, which suggests that improvements are possible at the level of the time-evolution method. A slight dependence of the system size on the onset of the breakdown further indicates that this might simply be a consequence of approaching a *volume law* state, see figure A.2.

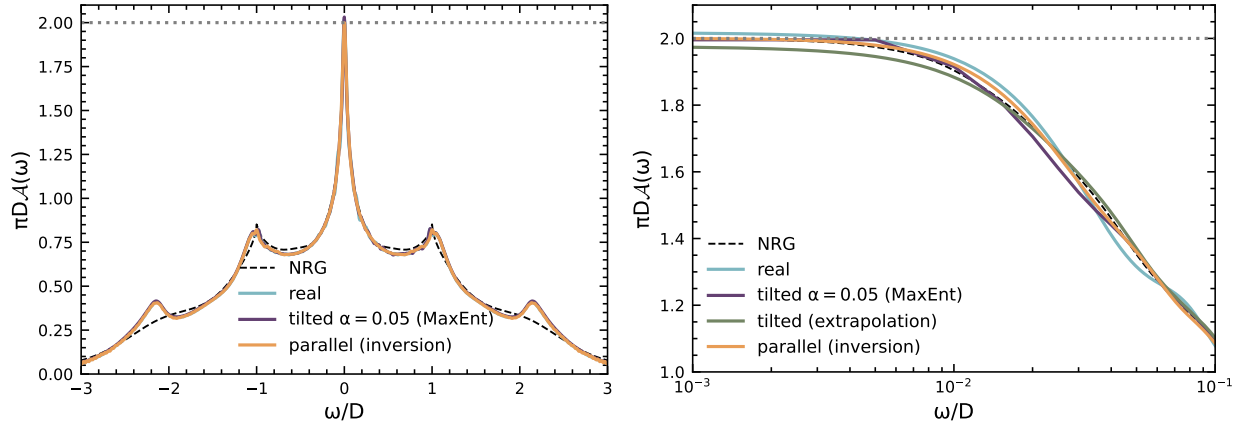


Figure 7.12: Spectral function $\mathcal{A}(\omega)$ for a three-orbital model with Hubbard–Kanamori interaction (left). All calculations are at truncated weight $w_t = 10^{-10}$ and without broadening ($\eta = 0$). Real-time results stem from a time evolution up to $t_{\max} = 120D^{-1}$ with a maximum bond dimension $m = 1024$. For the tilted (MaxEnt) method at angle $\alpha = 0.05$ the corresponding values are $t_{\max} = 180D^{-1}$ and $m = 2048$, for the parallel (inversion) method at $\tau = 1$, $t_{\max} = 220D^{-1}$ and $m = 1024$; in that case, an increased bath size $N_b = 139$ was used. The post-processing for the parallel (inversion) method used $n = 3$ contours at $\tau_k \in \{1, 1.15, 1.3\}$ with coefficients $a_k \in \{33.2, -57.8, 25.6\}$ (rounded). The Friedel sum rule is obeyed to an accuracy of 0.2% for the tilted (MaxEnt) data and 0.02% for the parallel (inversion) data versus 0.8% in the real-time calculation. (right) zoom into (left) with additional data obtained from the tilted (extrapolation) method using contours $\alpha \in [0.05, 0.1, 0.15, 0.2, 0.25, 0.3]$, averaging over all 4th order contributions. The oscillations in the real-time result are now clearly visible. The tilted (MaxEnt) result using $\alpha = 0.05$ shows a similar unphysical slight dip between $\omega/D = 10^{-2}$ and 10^{-1} as for the SIAM.

The instability is mended by continuing with a contour parallel to the real axis for larger times before instabilities occur, see figure 7.10. Figure 7.11 shows the excellent quality of the result for the spectral function of the SIAM. This method is at the moment the least performant in speed, roughly a speed-up of a factor three, but attractive due to the absence of post-processing. Methodological progress in the time-evolution methods may make it fully competitive.

7.4 Multi-Orbital Problems

The need for reliable high-performance calculations of low-frequency information is particularly pressing for multiorbital models used in realistic quantum embedding computations of strongly correlated materials [27, 29, 115], more specifically for transport computations. Accurate real frequencies quantum impurity solvers are rare, with the notable exception of NRG, which is, however, limited in the number of orbitals.




Figure 7.13: Imaginary part of the self-energy of the Hubbard–Kanamori Hamiltonian. The real-time result ("real") shows oscillations and fails at $\omega/D \approx 0.05$, missing the low-frequency physics. The dotted line indicates the ω^2 Fermi-liquid behavior of the self-energy. The tilted (MaxEnt) and tilted (extrapolation) methods do not reach this regime. The parallel (inversion) method (with $n = 2, 3$ with $\tau_k \in \{1.15, 1.3\}$ and $a_k \in \{8.7, -7.7\}$) is in very good agreement with NRG data but for a prefactor (see text) down to $\omega/D \approx 0.002$.

To benchmark our method in such a case, we use the three-orbital Anderson model with the Hubbard–Kanamori interaction, see equation (3.20). For the sake of consistency with real material calculations, we choose $U' = U - 2J$ [28]. Each impurity couples to a bath as in the SIAM. While it is not the most generic case (which would have non-diagonal bath couplings), we expect our approach to generalize without major difficulty. As for the density of states, we use the same semi-circular density of states as before and obtain the bath parameters as for the SIAM. Each bath is modeled by $N_b = 99$ bath sites if not stated otherwise. As interaction parameters, we consider $U/D = 2$ and $J/D = 0.3$. In all calculations, we set $\eta = 0$.

The simulations are carried out using 2TDVP until $t = 20D^{-1}$ and 1TDVP until the final time of $t = 180D^{-1}$. We use our novel MT3N tree-tensor network, which is tailored to multiorbital impurity models, see chapter 6. The TDVP implementation follows [199]. Again, we use NRG results as a benchmark since it is currently the most accurate method available for low frequencies. The Hubbard–Kanamori Hamiltonian in its band-degenerate form has an $SO(3)$ orbital symmetry which makes it accessible to standard multiorbital NRG [141, 142, 217, 218] (without the need for interleaving the Wilson chain [44, 218–221]). Note that we do not exploit the $SO(3)$ orbital symmetry in our MT3N calculations as it typically does not appear in multiorbital simulations of real materials.

Despite the tailored representation on a tree-tensor network, pure real-time calculations fail at high-frequency resolution. The complex-contour calculations, on the other hand, yield reliable results. In figure 7.12, we show that the spectral function results from the tilted (MaxEnt) method and from the parallel (inversion) method agree overall very well, whereas the real-time results show unphysical wiggles. There is, however, again a small discrepancy between the tilted (MaxEnt) and the parallel (inversion) methods at ω/D a bit above 0.01. In view of the results below, we interpret this as an inaccuracy of the tilted (MaxEnt) result. NRG results agree excellently at low frequencies but show the expected deviations for larger frequencies due to the logarithmic discretization of NRG.

We calculate the self-energy of the Hubbard–Kanamori Hamiltonian as in the case of the SIAM but add further methods to have some mutual benchmarking. We obtain $G(\omega)$ in four different ways: (i) from a real-time calculation; (ii) from a tilted (MaxEnt) calculation at $\alpha = 0.05$; (iii) from a tilted (extrapolation) calculation; and (iv) from a parallel (inversion) calculation at $\tau = 1.3$. In cases (ii) and (iv), we obtain $\mathcal{A}(\omega)$ and then $\text{Im} G(\omega)$ with the Kramers–Kronig transformation.

In Fig. 7.13, we observe that the real-time result shows weak but unphysical oscillations at frequencies above $\omega/D \approx 0.05$ and fully misses the low-energy Fermi-liquid physics $\text{Im} \Sigma(\omega) \sim \omega^2$ at lower frequencies. Both the tilted (MaxEnt) and tilted (extrapolation) methods perform somewhat better but also fail at low frequencies. At higher frequencies, the tilted (extrapolation) result perfectly agrees with the parallel (inversion) result.

The parallel (inversion) method easily reaches the ω^2 -regime. As in the case of the SIAM, sum rule violations in connection with using Dyson’s equation lead to a breakdown at very low frequencies. Interestingly, the breakdown does not occur at higher-frequency than in the SIAM case, indicating that this problem is not aggravated by the higher complexity of the studied model.

In the case of the three-band Hubbard–Kanamori model, NRG provides the best results for comparison but is no longer an *exact* benchmark: The large dimension of the local Hilbert space requires a rather large value of the discretization parameter, here $\Lambda = 6$ [43]. Averaging over n_z shifted discretization grids mitigates the effects of a coarse resolution of the hybridization function to some extent [123]. As discussed in [6], we found the best results by extrapolating to $n_z \rightarrow \infty$. Generally, we observe that the NRG result for the ω^2 coefficient of $-\text{Im} \Sigma$ decreases with increasing resolution of the hybridization function (decreasing Λ , increasing n_z). It is, therefore, remarkable that the MT3N result is slightly offset from the NRG result by a factor < 1 in the ω^2 regime before the breakdown occurs at $\omega/D \approx 0.002$.

We also considered the Dworin–Narath Hamiltonian, which can be represented similarly as in Eq. (3.20), except that one necessarily has $U' = U - J$ and that the last term in Eq. (3.20) (known as pair hopping) is missing [28]. It obeys a larger orbital symmetry than the Hubbard–Kanamori Hamiltonian (SU(3) instead of SO(3)), and thus allows for highly accurate NRG calculations at $\Lambda = 4$ (see [6]). Figure 7.14 shows the imaginary part




Figure 7.14: Imaginary part of the self-energy of the Dworin–Narath Hamiltonian. The parallel (inversion) method is in very good agreement with NRG data down to $\omega/D \approx 0.004$. All calculations are at a truncated weight $w_t = 10^{-10}$. Real-time evolution is up to $t_{\max} = 120D^{-1}$ with a maximum bond dimension $m = 512$. The tilted (MaxEnt) method uses an angle $\alpha = 0.05$ and goes to $t_{\max} = 180D^{-1}$ and maximum bond dimension $m = 1024$, the parallel (inversion) method uses $\tau = 1$ and goes to $t_{\max} = 140D^{-1}$ and maximum bond dimension $m = 1500$, the bath size was increased to $N_b = 139$. Results for the parallel (inversion) method are shown for $n = 4$ with $\tau_k \in \{0.85, 1, 1.15, 1.3\}$; the coefficients a_k (see main text) are then approximately $\{9.8, -14.2, 7.5, -2.1\}$.

of the self-energy for the Dworin–Narath Hamiltonian with the same parameters as before ($U/D = 2$ and $J/D = 0.3$). The performance of the time-evolution methods is similar to Fig. 7.13: going toward low frequencies, the real-time evolution soon yields unphysical results, the tilted (MaxEnt) and tilted (extrapolation) schemes improve on this, while the best result by far is obtained via the parallel (inversion) method. The corresponding curve follows the Fermi-liquid ω^2 behavior down to $\omega/D \approx 0.004$. Importantly, the agreement with NRG is excellent, as the coefficients of the ω^2 behavior match. This indicates that the small difference in the prefactor of the self-energy in Fig. 7.13 comes from the fact that the Hubbard–Kanamori NRG result is not fully converged in all numerical parameters.

The parallel (inversion) method emerges as the most reliable and performant among the methods tested here. Moreover, it can be further improved systematically: it rests on an analytically exact formula, and the breakdown at very low frequencies occurs because $G(t + i\tau)$ was not calculated for long enough times to yield highly reliable $G_\tau(\omega)$ for very small ω . As entanglement growth is curtailed, the region of ω^2 scaling can be extended to smaller frequencies by longer time evolutions. Note that the complex-time evolutions

employed here easily generalize to less symmetric situations, as we did not use any of the emerging larger symmetries and also have sufficient numerical efficiency to move to systems with more than three orbitals. Additionally, we expect that the accuracy of the self-energies provided by complex time evolution can be further improved by using improved estimators, as in NRG [179].

7.5 Summary

Complex-time evolution is an addition to the toolbox of tensor network simulations that offers high resolution when computing Green's functions at low frequencies at a fraction of the computational cost. At the same time, it maintains the quality of high-frequency data previously available. The speedup compared to real time evolution is highest for the tilted contour, about two orders of magnitude. MaxEnt, as a method of analytical continuation, provides very good, but not excellent agreement with real-time data where the latter is essentially exact.

The parallel contour calculations yielded the second highest speedup of (more than) an order of magnitude; post-processing by two different methods, inversion and extrapolation, provided excellent mutual agreement, allowing mutual control, as well as with benchmark data. Results were always more accurate than tilted (MaxEnt) results. In the very low-frequency regime for the self-energy of the three-band model, our most challenging calculation, the tilted (extrapolation) method, ran into difficulties, but the parallel (inversion) method was stable and accurate. The calculation of the self-energy of the Hubbard–Kanamori and Dworin–Narath three-band models reaches successfully into the $\text{Im } \Sigma(\omega) \sim \omega^2$ Fermi liquid regime down to $\omega/D \approx 0.002$ in agreement with NRG benchmark results. The parallel (inversion) method, therefore, seems to provide the best compromise in speedup and accuracy. It can be systematically improved to reach even lower frequencies, albeit at mounting numerical cost. Design decisions here will reflect compromises between required low-frequency resolution and available CPU time.

The availability of multiple complex time-evolution schemes that can be directly used or post-processed in different ways to extract real-time information makes these methods very controlled. We expect these methodological advances to be particularly useful in the context of quantum-embedding methods using frequency-based information, such as DMFT and its derivatives, where the efficiency of impurity solvers is paramount. In fact, the very high efficiency of the tilted contour calculations may make them the preferred approach for DMFT despite the limited accuracy. It remains subject to future research to apply these methods in the context of global quenches or higher-dimensional systems simulated directly on their real-space lattices. We expect that the suppression of high-energy contributions by imaginary time evolution may give access to long-time information inaccessible to real-time methods without substantial loss of accuracy.

Chapter 8

LiV₂O₄ - A Heavy Fermion TMO

The following results are based on the author's preprint [8].

The transition metal oxide lithium-divanadate (LiV₂O₄) generated significant interest after it was reported to exhibit heavy fermion behavior at low temperatures in 1997 by Kondo et al. [46]. Prior to this discovery, heavy fermion behavior had only been observed in rare earth elements, and it was essentially considered exclusive to f-orbital materials. However, LiV₂O₄ demonstrated the same characteristic drastic increase in quasiparticle mass at low temperatures, marking the first observed instance of such behavior in a d-orbital system. The significance of this extends beyond mere novelty; it necessitates a reevaluation of the foundational principles governing heavy fermion physics.

Previously, heavy fermion physics was understood to be intimately linked with the Kondo effect. In 1964, Jun Kondo [222] explained the high spectral densities near the Fermi edge and peculiar minima in the resistivity at low temperatures that had been reported for various alloys by considering the scattering of itinerate conduction electrons at localized magnetic moments. Essentially, this effect is not a direct consequence of specific band structures or crystal symmetries but of the presence of localized magnetic moments in the compound [222]. At high temperatures, conduction electrons scatter off these magnetic moments, while at low temperatures, they form a bound, non-magnetic state with the localized moments known as a Kondo singlet. Kondo, in his seminal work, already speculated that the strong correlation of itinerate electrons with magnetic moments would lead to a large enhancement of quasiparticle weight, foreshadowing the dawn of heavy fermion materials.

A decade after Kondo's findings, the first observations of compounds exhibiting large mass enhancements were reported. These early works were based on cerium alloys [223, 224], one of which was found to host superconductivity [224]. The latter study was the first observation of superconductivity in systems with strongly renormalized conduction electrons [224]. Moreover, studies by Mathur et al. revealed that some cerium-based alloys host unconventional, magnetically mediated superconductivity [225]. Overall, heavy fermion

systems host a myriad of functional properties [225, 226] that are a direct consequence of the strong electronic correlations within these compounds. Little is understood about these materials due to the immense complexity associated with their simulation. Thus, LiV_2O_4 is a rare opportunity to probe the mechanism behind heavy fermion physics and associated effects without the need to describe 4f-orbital alloys.

We will introduce LiV_2O_4 , detailing its structure and properties, and provide an overview of theories that aim to explain the emergence of heavy fermion physics in a d-orbital material in section 8.1. This will be followed by a discussion of our computational setup in section 8.2 and newly employed methods in section 8.3. Subsequently, we will present a detailed discussion of our results in section 8.4.

8.1 Material

LiV_2O_4 crystallizes in a face-centered cubic (fcc) spinel structure, identified by the space group $\text{Fd}\bar{3}\text{m}$, with crystallographic data obtained through neutron scattering experiments, as reported by Kondo et al. [46]. The primitive unit cell is composed of 14 atoms, including 4 vanadium (V) atoms. Each V atom sits in an octahedral crystal field exerted by the surrounding oxygen atoms. They form corner-sharing oriented octahedra with a trigonal D_{3d} local point group symmetry. The crystal field lowers the overall symmetry, breaking the t_{2g} degeneracy which results in a single a_{1g} orbital and energetically higher doubly degenerate e_g^π orbitals, see figure 3.2. The formal valency of vanadium is +3.5, which leads to an electron occupancy of 1.5 electrons per V site in this structure.

Despite extensive experimental and theoretical efforts aimed at understanding the mechanism behind the heavy fermion behavior observed in LiV_2O_4 , the underlying principles remain elusive. Nevertheless, we aim to provide a brief overview of previous findings and proposed theories regarding the emergence of heavy quasiparticles in this material. We will primarily focus on two prevalent theories. The first posits that the Kondo effect drives the emergence of heavy quasiparticles akin to 7-orbital transition elements. The second theory suggests that the large quasiparticle mass is a consequence of the geometric frustration inherent in the spinel crystal structure. Note that this idea is not based on a specific mechanism but rather a general concept. Geometric frustration leads to an increase in entropy in the material, which goes hand in hand with an increased specific heat coefficient, thus indicating the formation of heavy quasiparticles.

The work of Kondo et al. not only reported the formation of heavy quasiparticles at low temperatures but also firmly established the presence of local moments at higher temperatures. They observed a clear correspondence with the Curie-Weiss law up to the estimated Kondo temperature of roughly $T_K = 28$ K. However, no evidence for magnetic order was found above $T = 4.2$ K [46]. They interpreted their findings in the context of Kondo lattice models but also speculated, in accordance with [227], that the geometric frustration of the

crystal structure suppresses the onset of anti-ferromagnetic ordering and destabilizes the local moments at low temperature, ultimately leading to Fermi liquid (FL) behavior [46].

Anisimov et al. [228] proposed that the heavy fermion behavior can be understood by considering the interaction between local moments, i.e. electrons in the a_{1g} orbital and the spin of the itinerate electrons of the e_g^π orbitals. They argue that the strong Hund's coupling between these orbitals gives rise to an effective antiferromagnetic interaction between localized moments and spins, similar to the Kondo exchange. The Kondo theory was further substantiated by experimental findings by Shimoyamada et al. [163]. They observed a sharp quasiparticle peak right above the Fermi edge in their photoemission spectroscopy experiments that scales akin to a Kondo peak, i.e., both the width and position of this peak roughly scale as $\sim k_B T$. While this does corroborate the Kondo theory, it does not yet prove the presence of Kondo or Kondo-like correlations in the material.

Urano et al. found that, while the low energy behavior agrees well with that of a Kondo like material, the high energy transport behavior is substantially different from conventional f-orbital heavy fermion compounds [229]. Moreover, they argue that, structurally, LiV_2O_4 is close to a spin glass phase, yet again, indicating the importance of geometric frustration in this material [230]. Finally, orbital-resolved nuclear magnetic resonance measurements by Shimizu et al. [231] found no indication of Kondo physics down to low temperatures, signifying that the localized character of the a_{1g} orbital persists down to the Fermi liquid phase of LiV_2O_4 . In light of these findings, the theory of geometric frustration became significantly more popular and likely.

Lithium-divanadate exhibits a weak cusp in the magnetic susceptibility below $T = 20$ K [46]. However, no magnetic order has been observed down to $T = 0.02$ K, indicating strong frustration [46, 232]. Inelastic neutron scattering experiments of Tomiyasu et al. [232] indicate that the formation of heavy quasiparticles in LiV_2O_4 can be attributed to spin-orbit fluctuations. These fluctuations are argued to be a consequence of the formation of energetically degenerate magnetic states with a spatial correlation structure akin to antiferromagnetic polytetrahedra. DFT studies by Gong et al. [233] later confirmed the existence of ditetrahedron antiferromagnetic states with lower total energies compared to other magnetic configurations.

While most efforts towards LiV_2O_4 have focused on describing the emergence of heavy quasiparticles within the frameworks of Kondo physics or geometric frustration, it is important to highlight a third theory. The mechanism behind heavy fermion physics in d-orbital systems does not appear to be unique. Similar compounds, specifically CaRuO_3 [234] and Sr_2RuO_4 [235, 236], appear to host heavy quasiparticles due to their vicinity to either quantum critical points or metal-insulator transitions [237].

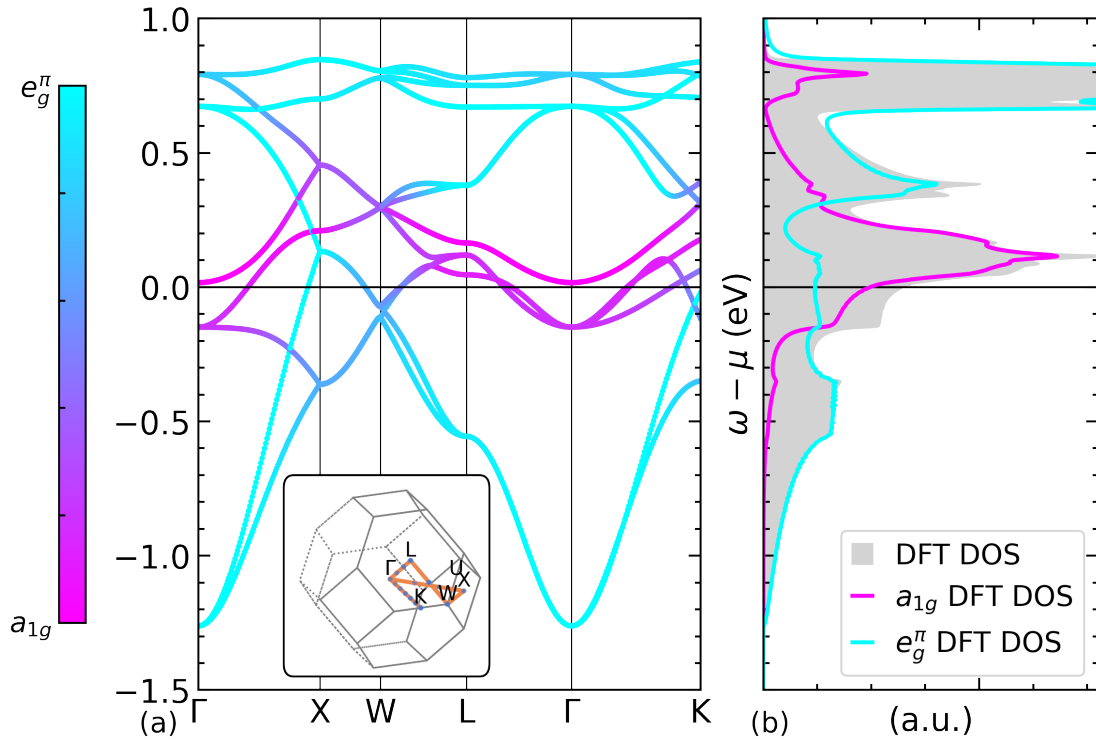


Figure 8.1: DFT band structure of LiV_2O_4 . Individual contributions of a_{1g} and e_g^π orbital are color-coded. The k-path between high symmetry points in the first Brillouin zone is shown in the inset. The k-summed density of states is shown in (right).

Arita et al. [238] proposed that the a_{1g} orbital in LiV_2O_4 can be modeled as a lightly doped Mott insulator and that heavy quasiparticles are a direct consequence of this. Their effective 2-orbital LDA+DMFT simulations (one a_{1g} orbital and one e_g^π orbital) revealed the emergence of a strong quasiparticle peak near the Fermi edge at low temperatures. They employed standard QMC techniques for a series of high-temperature spectral functions and complemented their data with $T = 0$ K results from their projective QMC method [239, 240]. The simulations showed an emerging quasiparticle peak near the Fermi edge, in accordance with experimental results [163]. This peak is preserved when the hybridization between the a_{1g} and e_g^π orbital is switched off [238]. They emphasized that this finding would directly contradict previous Kondo-like explanations based on the interaction between localized moments in the a_{1g} orbital and itinerant electrons in the e_g^π orbital. Moreover, they claim that the heavy fermion physics in lithium-divanadate is simply the result of the a_{1g} orbital being close to a Mott transition. Similarly, a recent work by Crispino et al. [241] argued that the emergence of heavy quasiparticles in d-orbital materials is the consequence of frustrated orbitally selective Mott transitions (OSMT).

8.2 DFT+DMFT

We performed DFT calculations using the QUANTUM ESPRESSO (QE) software package, where we fixed the structure parameters to match the experimental values reported for 12 K [46]. We construct maximally localized Wannier functions (MLWF) using WANNIER90 for the low-energy states near the Fermi level. These functions accurately replicate the QE bandstructure, as illustrated in figure 8.1. Integrating the projection of the density of states onto the a_{1g} and e_g^π orbital to the DFT Fermi level yields orbital occupations of 0.41 electrons per a_{1g} orbital per spin and 0.545 per e_g^π orbital per spin.

We construct a tight-binding Hamiltonian $H^{\text{W90}}(\mathcal{R})$ in real space from our MLWFs that accurately captures the essential low-energy physics of the compound. By employing Wannier interpolation, we Fourier transform $H^{\text{W90}}(\mathcal{R})$ onto a dense $41 \times 41 \times 41$ k -point mesh, thus avoiding any k -point discretization errors. This ensures high accuracy in our calculations, even at low temperatures. We apply this procedure to both an effective 12-band model (comprising a_{1g} and e_g^π orbitals) and an effective 20-band model (incorporating all five d-orbitals per V site). Additionally, we rotate the Hamiltonian into the crystal field basis on each of the 4 V sites to diagonalize the local non-interacting Hamiltonian and the hybridization function.

We performed constrained random phase approximation (cRPA) [127, 242] calculations to determine the effective Coulomb interaction within our low-energy model. By bipartitioning the electronic structure into a subspace near the Fermi level and the remainder of the system, we are able to compute the effective, partially screened Coulomb interaction in the static limit ($\mathcal{U}(\omega = 0)$), which serves as the effective interaction in our low-energy model. Additionally, we fit the full four-index effective interaction tensor to the three Kanamori parameters: U , U' , and J , see equation (3.20), to obtain a form that can be treated more efficiently by impurity solvers. The resulting Kanamori parameters are: $U = 3.94$ eV, $U' = 2.83$ eV, and $J = 0.56$ eV for the 3-orbital model and $U = 4.23$ eV, $U' = 3.13$ eV, and $J = 0.53$ eV for the full 5-orbital model.

DMFT calculations are performed using the TRIQS [243] software library, and its interface to electronic structure codes, TRIQS/DFTTOOLS [244], in a one-shot manner; that is, without updating the charge density in the DFT code. We solve the full interaction problem either using our tensor network based impurity solvers or the QMC solver in the hybridization expansion formulation, TRIQS/CTHYB[245]. All QMC based DMFT simulations are performed using TRIQS/SOLID_DMFT [246]. We consider the DMFT cycle to be converged when $|\mathcal{G}^{\text{loc}} - \mathcal{G}| < 10^{-4}$. Moreover, we closely monitor the convergence of the first few Matsubara frequencies in the self-energy to ensure that our results are fully converged, setting a threshold of ~ 1 meV for the change per iteration.

CTHyb calculations were performed from 290 K down to 11.6 K. Note that a significant increase in warm-up time, up to 1×10^5 warm-up cycles per Monte Carlo walker, was necessary to reliably ensure convergence. CTHyb results agree well with our TN solver down to 11.6 K, where the onset of heavy fermion behavior is already indicated. However, a further decrease of the simulation temperature is unfeasible for CTHyb due to the $\mathcal{O}(\beta^3)$ scaling of QMC algorithms.

We complement QMC results with our TN impurity solver, which offers a significantly higher energy resolution at very low temperatures. We extended our previously discussed TN-based solver, as detailed in chapter 4, to non-zero temperatures using thermal state purification [73, 247]. This enabled us to simulate temperatures as low as 2.9 K ($\beta = 4000 \text{ eV}^{-1}$). We utilized several time evolution techniques during the calculation of the impurity Green's function. Initially, the global subspace expansion method (GSE) [107] was employed until $\tau = 2 \text{ eV}^{-1}$, followed by two-site TDVP[90, 248] until $\tau = 25 \text{ eV}^{-1}$, and subsequently the local subspace expansion method (LSE) [7, 109] for the remainder of the time evolution. Time evolution was performed until $\tau = \frac{\beta}{2}$ for both zero temperature and finite temperature simulations. Additionally, zero temperature time-ordered Green's functions were extended using the linear prediction method [146] until they were fully decayed, as described in section 4.3.1.

We used a similar procedure for the preparation of our thermal states. Initially, we employed the global Krylov method [92] until $\tau = 0.5 \text{ eV}^{-1}$, followed by the global subspace expansion method (GSE) [107] until $\tau = 5 \text{ eV}^{-1}$, then the two-site TDVP[90, 248] until $\tau = 20 \text{ eV}^{-1}$, and finally the local subspace expansion method (LSE) [7, 109] until $\tau = \beta/2$. To limit the impact of truncation and projection errors we used exponentially increasing time steps $\delta\tau_n = \min(0.1 \times 1.01^n, 0.5) \text{ eV}^{-1}$ during thermal state preparation. Tensor network based time evolutions were performed with a truncated weight of $w_t = 10^{-11}$ and a maximum bond dimension of 1536 and 2048 $\text{SU}(2)$ states for 3 orbital and 5 orbital simulations, respectively. Ground state searches in our zero temperature calculations used a truncated weight of $w_t = 10^{-14}$ while allowing for a maximum bond dimension of 2048 and 4096 $\text{SU}(2)$ states for 3 orbital and 5 orbital simulations, respectively.

Significant deviations observed between self-energies obtained for $\beta = 1000$ by the QMC solver and $\beta_{\text{eff}} = 1000$ by the TN solver, required the extension of our TN framework to finite temperature simulations, see figure 8.2. This further corroborates that no direct correspondence between β and β_{eff} exists. The substantial deviation directly results from the low Fermi liquid temperature T_{FL} in LiV_2O_4 . Our finite temperature simulations at $\beta = 1000$ correspond to an absolute temperature of 11.6 K, while T_{FL} is suspected to be below 10 K [249]. Thus, we only find a mere onset of a potential Fermi liquid behavior in finite temperature simulations, while zero temperature simulations already show a more pronounced emergence of a Fermi liquid phase.




Figure 8.2: Comparison between self-energies obtained via TN or QMC methods at $\beta = 1000$ and $\beta_{\text{eff}} = 1000$. The a_{1g} orbital (left) shows a significant slope at low frequencies, indicating the large mass enhancement characteristic of this compound. However, we observe a stark difference between zero and finite temperature results. This discrepancy directly results from the low Fermi liquid temperature in this material [249]. Only the zero temperature self-energy shows a clear indication for the formation of a Fermi liquid. However, we find a good agreement between both methods at $\beta = 1000$.

However, we find a good agreement between the QMC and finite temperature TN results overall. The presence of small deviations at low energies will be discussed in section 8.3.2. This finding is a remarkable testament to the efficiency and maturity of both approaches, as both yield highly precise results even in largely unexplored territory.

8.3 Purification

Tensor network-based solvers have traditionally been used primarily for zero-temperature simulations. However, accurate comparisons with QMC and, more importantly, experimental results require finite-temperature data. Little is known about finite temperature TN simulations in the context of DMFT. Recently, Cao et al. [250] proposed the use of minimally entangled typical thermal state (METTS) for TN based impurity solvers. Yet, the statistical averaging required in this approach leads to a significantly increased computational cost compared to zero temperature simulations. Thus, we propose a conceptually simpler method for the simulation of materials at finite temperatures: purification. We refer to [73, 247] for a detailed review on how to combine purification with tensor network methods.




Figure 8.3: Depiction of an infinite temperature MPS in the purification framework. Each physical site (blue) has a neighboring ancilla counterpart (orange). Thus, the overall Hilbert space size is doubled. These pairs are mutually maximally entangled, resulting in an alternating bond dimension pattern.

Any mixed state in physical space can be represented as a partial trace over a Schmidt decomposition of a pure state [73]:

$$\hat{\rho}_P = \sum_{i=1}^m s_i |i\rangle_P \langle i|_P = \text{Tr}_A |\psi\rangle \langle \psi|, \quad (8.1)$$

$$|\psi\rangle = \sum_{i=1}^m s_i |i\rangle_P |i\rangle_A, \quad (8.2)$$

where s are singular values, P denotes our physical Hilbert space, and A is an ancilla space, which we assume to be identical to P . In MPS representation, this corresponds to adding a neighboring ancilla site to each physical site, thus doubling the overall Hilbert space dimension. See figure 8.3 for a depiction of a MPS in the purification framework. Inserting an identity into the density matrix of a thermal state ρ_β allows us to rewrite it as

$$\hat{\rho}_\beta = \frac{1}{Z(\beta)} e^{-\beta \hat{H}} = \frac{1}{Z(\beta)} e^{-\beta/2 \hat{H}} \hat{I} e^{-\beta/2 \hat{H}}. \quad (8.3)$$

Identifying $\hat{I} = \hat{\rho}_{\beta=0} Z(\beta=0)$ and expressing $\hat{\rho}_{\beta=0}$ as a purification allows us to rewrite above equation as

$$\hat{\rho}_\beta = \frac{Z(0)}{Z(\beta)} e^{-\beta/2 \hat{H}} \text{Tr}_A |\psi_0\rangle \langle \psi_0| e^{-\beta/2 \hat{H}}. \quad (8.4)$$

As \hat{H} does not act on A we can pull out the trace, which finally leaves us with

$$|\psi_\beta\rangle = e^{-\beta/2 \hat{H}} |\psi_0\rangle. \quad (8.5)$$

Thus, to obtain a finite temperature state, we must first find the infinite temperature state $|\psi_0\rangle$. However, as $\hat{\rho}_0$ factorises as

$$\hat{\rho}_0 = \frac{1}{d^L} \hat{I} = \left(\frac{1}{d} \hat{I}\right)^{\otimes L}, \quad (8.6)$$

we can simply purify the local mixed state on each physical site i as [73]

$$|\psi_{i,\beta=0}\rangle = \sum_{\sigma} \frac{1}{\sqrt{d}} |\sigma\rangle_P |\sigma\rangle_A. \quad (8.7)$$

Hence, the infinite temperature state can be easily constructed from maximally entangled states between physical sites and their ancilla counterpart. This results in an alternating bond dimension pattern in MPS as sites are only pairwise entangled; see figure 8.3. There is a certain level of freedom in the choice of basis for these states. We strongly advise constructing the state so that it conserves good quantum numbers of the Hamiltonian [198]. For electronic systems, we construct our maximally entangled state as

$$|\psi_{i,\beta=0}\rangle = \frac{1}{\sqrt{4}} \left(|\uparrow_P \downarrow_P\rangle + |\uparrow_A \downarrow_P\rangle + |\uparrow_P \downarrow_A\rangle + |\uparrow_A \downarrow_A\rangle \right). \quad (8.8)$$

This conserves the particle number N , total spin S , and spin S^z throughout time evolution. Note that this still allows for a grand canonical state after tracing out the ancilla space A . For $SU(2)$ invariant systems, as is the case here, one can construct the maximally entangled state as

$$|\psi_{i,\beta=0}\rangle = \frac{1}{\sqrt{3}} \left(|0_P 2_A\rangle + |1_P 1_A\rangle + |2_P 0_A\rangle \right), \quad (8.9)$$

where $\{0, 1, 2\}$ denote the number of particles on each site.

8.3.1 How to stabilize finite temperature simulations

While we are able to simulate unprecedented low temperatures in real material simulations, we are once again confronted with the presence of a gap in the time-ordered Green's function. Refer to section 4.3 for a discussion on how we addressed this issue in zero-temperature simulations. We observe that G^p and G^h do not perfectly coincide at $\tau = \beta/2$, this time, however, they must. Moreover, our previous solution of extending $G(\tau)$ via linear prediction is no longer applicable due to the periodicity and the therein implied growth near $G(\tau = \beta/2)$.

The presence of a gap in the time-ordered Green's function is not merely the consequence of emerging inaccuracies in challenging material simulations. Even relatively simple simulations of the Hubbard model at $U = 2D$ at $\beta = 8$ with a semi-elliptical density of states reveal a gap at $G(\tau = \beta/2)$, see figure 8.4. Moreover, the general behavior of the gap with system size is inconsistent. While simulations with a single bath site at half-filling show no gap, we observe the largest gap away from half-filling, see figure 8.4. The presence of the gap was neither impacted by varying the timestep not using using high-precision global time evolution methods. Yet, around 5×10^{-4} , the gap stabilizes, both for the simple test case and for LiV_2O_4 at very low temperatures.

In our zero-temperature simulations, we relied on linear prediction to mitigate the effects of the gap. Specifically, we encountered two issues: the gap at $G(\tau = \beta_{\text{eff}}/2)$ and an improper description of low-energy physics. We demonstrated that a simple shift in the time particle/hole Green's function mitigates all observed oscillations in the real part of the Green's function while leaving the imaginary part, which shows no oscillations, unchanged, as illustrated in figure 4.3.




Figure 8.4: Comparison of gaps in the time-ordered Green's function at $\tau = \beta/2$. We show the evolution of the gap in a one-band Hubbard model with $U = 2D$ and semi-elliptical density of states at $\beta = 8$ in (left). While simulations at half-filling ($\mu = 0$) initially display no gap, a clear formation of a gap is observed with larger bath sizes. The origin of this gap remains somewhat mysterious. We find a decrease in the gap with bath length for systems away from half-filling. This ultimately suggests that the gap is not purely a consequence of inaccuracies in our simulations. Moreover, the use of high-precision global time evolution methods does not impact the presence of this gap. We illustrate the evolution of this gap in (right) over a wide temperature range for LiV_2O_4 . The absolute size of the gap remains stable for all considered temperatures.

Moreover, due to the periodicity of the time-ordered Green's function, we do not expect errors arising from an insufficiently long time evolution. Therefore, we can reasonably assume that a shift, as defined in equation (4.11), effectively mitigates all effects of the gap.

We showcase the effectiveness of the shift on the example of LiV_2O_4 at $\beta = 1000$. We find that oscillations in the real part of the impurity Green's function are completely mitigated, see figure 8.5. Upon careful inspection, we find that the maximum of the time-ordered Green's function is not located at $\tau = \beta/2$ but rather slightly to the left of it, again, hinting at a possible minor source of error. Resampling $G(\tau)$ on the dlr grid [164, 165] shifts the maximum ever so slightly towards the middle. Although the impact of this approach is seemingly minor, it significantly improves the high-energy behaviour and, by that, the tail fitting accuracy of our Green's function [34, 251]. The shift itself is sufficient to unveil the onset of the heavy-fermion behaviour in the self-energy as indicated by the stark trend of its imaginary part towards zero, see figure 8.6. Moreover, the progressively smaller amplitudes of the oscillations in the self-energy at lower temperatures corroborate the accuracy of the reconstruction of the proper self-energy via the shift.




Figure 8.5: Presence of a minute gap at $\tau = \beta/2$ in our simulations of the a_{1g} orbital at $\beta = 1000$, (left). We can mitigate oscillations in the impurity Green's function (right) that arise as a consequence of this gap by shifting the particle and hole Green's functions towards each other. This approach leaves the imaginary part of $G(i\omega_n)$ invariant. However, other minor sources of error, for instance, a slightly shifted position of the maximum in $G(\tau)$ persist.




Figure 8.6: Shifting G^p and G^h towards each other reveals the heavy fermion behaviour in the a_{1g} orbital at $\beta = 1000$ (left). Moreover, the accuracy of the reconstruction is corroborated by the ever-smaller oscillations (right) with decreasing temperature in the original data. Minor deviations in the self-energy arise when combining the shift with the DLR grid resampling, likely due to the improved high-energy behavior of the resampled Green's function.

8.3.2 Odd/Even effect

Tensor network based impurity solvers do not work directly in the thermodynamic limit, unlike QMC solvers. Despite this seemingly fundamental difference, they have often been found to be in good agreement with each other [1, 4]. However, this agreement is partly due to these studies' comparably low energy resolution. The low temperatures accessed by both solvers in the context of this work reveal a more concerning picture regarding finite-size effects in tensor network impurity solvers.

We observed significant odd/even effects in the number of bath sites, both in the zero temperature as well as the finite temperature results, albeit in a weaker form. We find that *even* bath lengths result in a zero temperature ground state sector with an *even* total number of particles and total spin $S = 0$. Conversely, *odd* bath lengths favor a ground state with an *odd* total number of particles and total spin $S = 0.5$. This changes the overall band parity of the a_{1g} band, thus significantly influencing its low-energy behavior. Consequently, this affects the observed onset of the heavy fermion behavior.

The emergence of a heavy Fermi liquid phase in this compound is intimately linked with the disappearance of the local magnetic moment in the a_{1g} orbital. The presence of the local moment results in a strong scattering amplitude, as shown in figure 8.2. The magnetic moment diminishes with decreasing temperature, ultimately resulting in the emergence of Fermi liquid behavior and the therein implied $S = 0$ state. We find a stark deviation between results with an *odd* or *even* number of bath sites, see figure 8.7. For *even* numbers of bath sites we report a clear onset of Fermi liquid physics, while *odd* bath lengths show no such indications. In contrast, the e_g^π orbital shows no noteworthy dependence on the number of bath sites. This highlights that the pronounced *odd/even* effects are a consequence of the strong correlation effects and the low Fermi liquid temperature in this material.

Some of the on-site energies in the bath of the a_{1g} orbital, determined during our discretization procedure, are $\sim \mathcal{O}(10^{-4})$. This results in the observed stark volatility of the total parity in the a_{1g} band. We want to showcase that the strong *odd/even* effects are merely a consequence of this instability of the ground state quantum number sector rather than the result of a fundamentally different hybridization function.

As the small on-site energies imply a small energy gap to the lowest eigenstate in the $S = 0$ quantum number sector, we can simply select the $S = 0$ state as our lowest energy eigenstate in zero temperature simulations without encountering severe norm instabilities during time evolution. Calculations at $\beta_{\text{eff}} = 1000$ show a clear onset of heavy fermion physics, even for *odd* bath sizes, see figure 8.8. Thus, we conclude that this behaviour is independent of the number of bath sites, but strongly reliant on an $S = 0$ ground state. Given that we would expect to find an $S = 0$ ground state in the FL phase, this aligns perfectly with our physical expectation.




Figure 8.7: Impact of the *odd/even* effect on LiV_2O_4 orbitals. We find that the emergence of the heavy fermion behavior in the a_{1g} orbital (left) at $\beta = 1000$ is critically dependent on an *even* number of bath sites L_b , which favors a total spin of $S = 0$ in zero temperature simulations. The local moment in the a_{1g} orbital still prevails for *odd* bath lengths, i.e., no indication for the onset of heavy Fermi liquid behavior can be found. In contrast, we observe little to no effect on the e_g^π orbital (right), corroborating that the presence of these strong *odd/even* effects is intimately linked with the presence of magnetic moments and low Fermi liquid temperature scale in this material.

We consider *even* bath sizes as the correct physical choice for LiV_2O_4 . This choice is further justified by the experimentally observed emergence of Fermi liquid behavior at very low temperatures. Additionally, it is corroborated by the consistency with QMC results. This comparison, see figure 8.2, reveals an overall good agreement between self-energies obtained by both methods. Upon closer inspection, one finds that QMC results exhibit a slightly smaller uptick at low frequencies compared to our MPS-based results. We consider this small deviation to be a potential consequence of the $S = 0$ bias that our choice of bath size may have introduced into the model. We consider such stark *odd/even* effects to be a rare consequence of the interplay between local magnetic moments and the low coherence temperature in this material. Given the overall stability of the e_g^π orbital self-energy for all considered bath lengths, we remain confident in the accuracy of the tensor network solver.




Figure 8.8: Self-energy in the a_{1g} orbital at $\beta_{\text{eff}} = 1000$ for *odd* bath lengths from the $S = 0.5$ ground state and the energetically closest $S = 0$ state. We find that the emergence of heavy fermion physics at this fictitious temperature is independent of the bath size and only depends on the total spin.

8.4 Results

The heavy Fermi liquid regime in LiV_2O_4 has so far been inaccessible to DMFT calculations due to the very low FL coherence scale [46, 238]. However, our recent advances in tensor network based impurity solvers allow us to access the low temperatures necessary to describe the emerging heavy quasiparticle phase. We complement these results with QMC data obtained at higher temperatures to study the full crossover from the incoherent metal to the low-temperature heavy FL regime.

Results are shown for the a_{1g} and e_g^π orbital of our 3-orbital model. The small electronic occupation of the e_g orbitals, $n_{e_g} = 0.01$, does not warrant a complete treatment on the basis of the full 5-orbital simulations. Nonetheless, we will present a comparison of 3-orbital and 5-orbital self-energies in appendix B.




Figure 8.9: Onset of heavy Fermi liquid behaviour in the a_{1g} orbital (left). This is marked by a steep trend of the self-energy towards zero for temperatures below approximately 10 K, in agreement with experimental predictions [249]. The a_{1g} and e_g^π orbital are denoted by circles and crosses, respectively. We firmly establish the mutual agreement of our TN and QMC impurity solvers over a wide range of temperatures. Additionally, we show the development of the sharp quasiparticle peak near the Fermi edge in the a_{1g} orbital (right) with decreasing temperature, as has been observed experimentally in [163].

We establish a good agreement between TN and QMC self-energies over a wide range of temperatures, see figure 8.9. Our tensor network results reveal the emerging heavy FL regime below 10 K, which coincides well with experimental predictions [249]. The a_{1g} self-energy can be described in two regimes. The high-temperature regime above $\omega_n \approx 0.006$ (eV) is governed by the formation of a local magnetic moment, as shown in figure 8.12. This leads to a strong scattering rate $-\text{Im}\Sigma_{a_{1g}}(i0^+)$, as can be inferred from figure 8.9. The scattering rate drops dramatically in the low-temperature regime below $\omega_n \approx 0.006$ (eV), indicating the emergence of a heavy Fermi liquid phase. The spectral function shows the formation of a dominant, temperature-dependent quasiparticle peak near the Fermi edge, as observed experimentally [163]. In contrast, the e_g^π self-energy depends only weakly on temperature.

The direct correspondence between the inverse quasiparticle renormalization $Z^{-1} = 1 - \frac{\partial \text{Im} \Sigma}{\partial i\omega_n} \Big|_{(i\omega_n \rightarrow 0)}$ and the mass enhancement $m^*/m = Z^{-1}$ within our DMFT approximation breaks down for multi-orbital models [1]. We thus consider the specific heat coefficient given as:

$$\gamma = \frac{\pi^2 k_B^2}{3} \sum_{m\sigma} A_{m\sigma}(0)/Z_{m\sigma}, \quad (8.10)$$

where m and σ denote orbital- and spin indices, respectively.




Figure 8.10: Self-energy for various Hund's coupling strengths at $\beta = 1000$ (left), other parameters are kept as is. We see that the onset of Fermi liquid behavior is progressively pushed to smaller frequencies and, in that sense, temperatures. We show the mass enhancement and occupation of the a_{1g} orbital for all considered J . We see that the onset of heavy fermion physics is clearly driven by Hund's coupling. Upon reaching an orbital occupation of $\approx n = 0.9$, we find a consistently high mass enhancement that no longer depends strongly on J . Note that this is only visible for $\beta_{\text{eff}} = 1000$. Finite temperature simulations are still well above T_{FL} for larger J . Thus, the mass enhancement appears to be shrinking.

This accurately measures the mass enhancement in multi-orbital systems, which is directly comparable with experimental results. We observe a rapid increase in γ in the low-temperature regime. At $T = 11.6$ K we have $\gamma = 159.24$ mJ/mol K^{-2} , whereas at $T = 2.9$ K, we have $\gamma = 3552.52$ mJ/mol K^{-2} . Additionally, we obtain a good agreement between our estimates of the specific heat coefficient within our DFT model, $\gamma_{\text{DFT}} = 17.09$ mJ/mol K^{-2} , and previous estimates [252, 253]. Values are provided per formula unit of LiV_2O_4 . Note that in the high-temperature regime of our DMFT model, we observe only a small, near-constant increase $\gamma/\gamma_{\text{DFT}} \approx 1.4$, indicating that the material is only weakly correlated. However, the strong enhancement in the low-temperature regime clearly necessitates our DFT+DMFT approach to accurately capture the intricate correlation effects.

The formation of a strong quasiparticle peak was also reported in a previous, approximate DMFT study by Arita et al. [240]. They proposed that LiV_2O_4 behaves akin to a weakly under-doped Mott insulator (1-2% doping). Moreover, they classified the e_g^π orbitals as mere spectators reducing their role to that of electron donors. Indeed, the orbital occupations of the DFT calculations read $n_{a_{1g}} = 0.41$ and $n_{e_g^\pi} = 0.545$, which differ significantly from the fillings obtained in our DMFT calculations of roughly $n_{a_{1g}} = 0.9$ and $n_{e_g^\pi} = 0.3$.

While this is in qualitative agreement with the mechanism proposed by Arita et al., the filling of the a_{1g} orbital is much lower than the necessary 1-2% doping that would result in a similar mass enhancement in a simple 1-band model. Moreover, we find that Hund's coupling is essential for the formation of heavy quasiparticles in this material. Hund's coupling itself plays multiple key roles, which we shall establish below. It promotes the significant redistribution of electrons from the e_g^π to the a_{1g} orbital. In the absence of Hund's coupling, the orbital occupation in our DMFT study is roughly $n_{a_{1g}} = 0.66$ and $n_{e_g^\pi} = 0.42$. More importantly, we find no heavy quasiparticles at $\beta = 1000$, see figure 8.10. In fact, we find that the heavy fermion behavior is critically linked to the occupation of the a_{1g} orbital, which emerges upon reaching an occupation close to $n_{a_{1g}} \approx 0.9$.

We find a maximum in the inverse quasiparticle renormalization around $J = 0.3$ for $\beta = 1000$ and $J = 0.4$ for $\beta_{\text{eff}} = 1000$. We expect that this is not a true maximum and that it will shift to slightly larger Hund's coupling strengths upon lowering the temperature. We do not consider the observed, drastic decrease in Z^{-1} for larger J to be physically relevant. It is purely an effect of the chosen temperature, as the comparison with zero temperature data indicates, see figure 8.10. We see that the increase in quasiparticle mass is strongly correlated with the increase in the occupation of the a_{1g} orbital, both of which are driven by Hund's coupling, reinforcing its critical role in this material.

While the description of LiV_2O_4 as an effective single orbital model close to a Mott transition appears far-fetched, the general importance of the orbital occupation cannot be understated, also in the presence of Hund's coupling. The total filling of LiV_2O_4 positions it right at the edge of the spin freezing regime of the Hubbard Kanamori Hamiltonian [254]. This regime is characterized by the formation of slowly decaying local magnetic moments. These local moments lead to the creation of heavy quasiparticles, driven by the spin-blocking mechanism in Hund's metals, where electrons can only transition between configurations with maximum spin.

So far, this regime has, to the best of our knowledge, only been firmly established within degenerate model systems, far astray from our context here. However, as we will illustrate below, the strong renormalization effects survive the presence of a crystal field splitting and could serve as a possible explanation for the emergence of heavy quasiparticles.




Figure 8.11: Transition towards the spin freezing regime of the Hubbard-Kanamori Hamiltonian with increasing orbital occupation n . Simulations have been performed at $\beta = 1000$ using the CT-Hyb impurity solver for $U/D = 3.94$ and $J/D = 0.56$. We observe the emergence of the characteristic up-tick in the self-energy indicating the transition from the bad metal to the Fermi liquid regime. For comparison, we show $\Sigma_{a_{1g}}$ in gray. The inset shows the inverse quasiparticle renormalization for different fillings. We find that characteristics of the spin freezing regime remain stable upon the introduction of a crystal field splitting (right). This stability indicates that the spin-blocking mechanism is a possible cause behind the emergence of heavy quasiparticles in LiV_2O_4 .

We consider the Hubbard-Kanamori model on the degenerate Bethe lattice with interaction parameters corresponding to the a_{1g} orbital of LiV_2O_4 , i.e., $U/D = 3.94$ and $J/D = 0.56$. Simulations were performed at $\beta = 1000$ using the CTHyb solver. We find a strong increase in Z^{-1} upon approaching half-filling, see inset figure 8.11. Moreover, we see the same characteristic strong uptick of the self-energy that we observe in LiV_2O_4 . Note that the Fermi liquid temperature in this model system is still significantly higher compared to LiV_2O_4 , as indicated by the earlier onset of the kink in the self-energy, see figure 8.11. These results are in line with previous works on the spin freezing regime in Hund's metals [254, 255].

The stark increase in quasiparticle mass at low temperatures even prevails upon the introduction of a crystal field splitting. Note that the dominant kink at small frequencies vanishes; however, the inverse quasiparticle renormalization remains large. The comparison with the fully degenerate model corroborates that the strong correlation effects do not primarily originate from the vicinity of the electron occupation to half-filling but rather the vicinity to an orbitally selective Mott transition (OSMT). In LiV_2O_4 , the a_{1g} orbital is similarly close to an OSMT.




Figure 8.12: Local spin-spin susceptibility in LiV_2O_4 . The spin-spin correlation in the a_{1g} orbital is nearly constant in τ leading to a clear temperature dependence in $\chi_m^{S_z S_z}(\nu_n = 0, T)$ (right). We fit the temperature dependence with the Curie-Weiss law $X = \frac{C}{T - \Theta}$ and find a near-perfect correspondence. However, no evidence for magnetic order above 4.2 K has been observed experimentally [46]. In comparison, the spin-spin correlation in the e_g^π orbital decays rapidly, showing no trace of a magnetic moment.

A further hallmark of the spin-blocking mechanism is the presence of a strong local magnetic moment, persisting down to very low temperatures. We want to establish the presence of such a local magnetic moment by considering the local spin-spin susceptibility $\chi_m^{S_z S_z}(i\nu_n)$, where ν_n denotes the bosonic Matsubara frequencies. We can compute this quantity via Fourier transformation of the local spin-spin correlation function in imaginary time:

$$\chi_m^{S_z S_z}(i\nu_n) = \int_0^\beta d\tau e^{i\nu_n \tau} \langle \psi_0 | \hat{S}_{z,m}(\tau) \hat{S}_{z,m} | \psi_0 \rangle, \quad (8.11)$$

where \hat{S}_z is a spin operator and m runs over all orbitals.

We find evidence of local magnetic moments in the a_{1g} orbital, that persist down to very low temperatures. This is indicated by the slow decay of $\chi_m^{S_z S_z}(i\nu_n)$, as shown in figure 8.12. The spin-spin correlation function becomes nearly flat at the considered temperatures, leading to a spin-spin susceptibility $\chi_m^{S_z S_z}(i\nu_n = 0, T)$ that is proportional to β , indicating that the Curie-Weiss law holds [255]. Indeed, we find a nearly perfect fit between the Curie-Weiss law and our simulations, see figure 8.12 (right).

Experiments showed that the spin-spin susceptibility in LiV_2O_4 follows a Curie-Weiss law down to 16 K where it experiences a cusp and becomes nearly temperature independent afterward [46]. Yet, we do not find a similar cusp in our simulations. This breakdown of the Curie-Weiss law reported in experiment is expected at the crossover into the Fermi

liquid regime. However, the behaviour in our model indicates a lower T_{FL} in our model compared to values reported in experiment.

While the spin-blocking mechanism appears compelling, the large filling of the a_{1g} orbital makes it cumbersome to disentangle it from other correlation effects that might arise in the vicinity of a Mott transition [221]. Recent studies even suggested frustrated orbitally selective Mott transitions (OSMT) as the mechanism behind the emergence of heavy fermion physics in d-orbital materials [241]. However, given that the a_{1g} orbital is still relatively far removed from an actual OSMT, we consider the spin-blocking mechanism of Hund's metals to be the main contributor to the emergence of heavy quasiparticles in LiV_2O_4 .

8.5 Summary

We leveraged recent advances in our tensor network impurity solver to reach unprecedentedly low temperatures in the simulation of materials using DFT+DMFT studies, incorporating the full Hubbard-Kanamori interaction. This allowed us to probe the emergence of heavy quasiparticles in the a_{1g} orbital of the transition metal oxide LiV_2O_4 .

Previous theoretical studies attributed the heavy fermion characteristic of this compound to its geometric frustration [230, 231] and mechanisms akin to the Kondo effect [228]. A more recent study by Arita et al. proposed that this material can essentially be understood as a single orbital model close to a Mott transition.

Contrary to the claims by Arita et al., we find that Hund's coupling plays a fundamental role in the formation of heavy quasiparticles in LiV_2O_4 . It drastically shifts the electron occupation towards the a_{1g} orbital, increasing its filling from $n_{a_{1g}} = 0.66$ to roughly $n_{a_{1g}} = 0.9$. This increase in orbital electron occupation is closely associated with the increase in mass enhancement.

We show that characteristics of the spin-freezing regime of the Hubbard-Kanamori Hamiltonian persist in the presence of a crystal field splitting. This makes spin-blocking a viable mechanism behind the intricate low energy physics of LiV_2O_4 . We demonstrated the existence of local magnetic moments that persist to very low temperatures. However, it remains unclear whether this might be the dominant effect behind the formation of heavy quasiparticles in d-orbital materials.

Further investigations are needed to disentangle the apparent similarity with the spin freezing regime and the strong correlation effects that might arise through a frustrated orbitally selective Mott transition of the a_{1g} orbital. It is important to note that both of these effects are largely independent of the band structure. Therefore, we consider it crucial to identify additional transition metal oxides that exhibit similarly low-energy physics as LiV_2O_4 to better understand these intertwined mechanisms.

Chapter 9

Conclusion and Outlook

We presented methodological advances in the simulation of strongly correlated materials using dynamical mean field theory and a study of the low-energy physics of the transition metal oxide LiV_2O_4 . Our recent advances enabled us to directly study the emergence of heavy fermion physics in LiV_2O_4 , allowing us to propose a new theory behind the formation of heavy quasiparticles in this d-orbital compound.

We introduced a new tensor network structure, the MT3N, that is tailored towards impurity models. This development generated a 10x speed-up for 5-orbital models in our real time impurity solver, enabling us to perform intricate real material simulations directly on the real frequency axis. The imaginary time impurity solver also benefits from the more efficient representation of correlation in the MT3N. Additionally, we generated an impressive 3x runtime improvement by the usage of recently developed subspace expansion algorithms during time evolution. These subspace expansion techniques proved crucial in extending this solver to finite temperatures. Our purification approach allowed us to reach temperatures as low as $T = 2.9$ K, an unprecedentedly low temperature for DFT+DMFT material studies incorporating the full Hubbard-Kanamori interaction.

Despite these advances, accurately describing real materials across the full frequency range remains challenging. To address this, we introduced a new analytic continuation algorithm, MinKL, that combines results obtained on the Matsubara axis from our imaginary time solver and real frequency data from our real time solver. This significantly improved the stability of the otherwise ill-conditioned analytic continuation procedure. Moreover, MinKL preserves the highly accurate low-energy description from our Matsubara axis data and faithfully extends it with high-energy information from the real time solver results. This allows us to obtain highly accurate real frequency data, even for challenging multi-orbital material simulations, enabling us to compare our numerical analysis with experimental observations.

Lastly, we introduced a novel time evolution concept by evolving systems along complex time contours [6]. By shifting time evolution away from the real time axis, we significantly curtail entanglement growth, enabling substantial improvements in accuracy and efficiency. This is achieved by systematically suppressing high-energy states via imaginary time evolution. We discussed several complex time contours along with multiple post-processing methods that analytically continue our results back to the real frequency axis. Among these, we introduced a tilted contour that posts a staggering 100x speed-up compared to real time contours. However, this tremendous gain in efficiency comes at a loss of accuracy due to the necessity of an ill-conditioned analytic continuation. Most notably, we presented contours that either require no post-processing or where the analytic continuation simplifies to a mere multiplication. The latter, parallel (inversion), method pars the precision of NRG down to $\omega/D \approx 0.002$, a level of precision previously out of reach for tensor network based real time impurity solvers all while preserving an impressive speed-up of more than an order of magnitude.

Outlook

While we presented substantial advances in our tensor network impurity solvers, we still see room for improvement. Although the imaginary time solver is not plagued by entanglement growth during time evolution, the need for increasingly longer time evolutions upon decreasing temperatures is cumbersome. Therefore, it is worthwhile to explore the option of directly computing the Matsubara Green's function in frequency space. This could be a promising endeavor given the likely low number of poles needed to accurately describe it.

The vast number of bath sites needed to accurately capture low-energy physics in the real time solver creates efficiency bottlenecks in otherwise insignificant parts of tensor network simulations. More precisely, creating both the Hamiltonian and a good initial state consumes considerable parts of the overall runtime of a real time DMFT iteration. We thus urge replacing the generic TNO construction [256] in favour of a finite state machine based construction principle [257]. Additionally, the importance of filesystem performance must not be underestimated. At times, we spent nearly half of our runtime on I/O operations in our imaginary time solver.

Complex time evolution has the potential to solve the basis optimization problem in real time evolution. While basis transformations like the natural orbital basis [41] have shown impressive speed-ups in DMRG, they fell short during time evolution. This is unsurprising as these transformations are typically not dynamically adapted to the evolving state [195]. As we remain close to an eigenstate during complex time evolution, it is plausible that the optimal basis also remains similar throughout time evolution.

There is still much to explore in the context of complex time evolution. It is yet unclear how to best close the self-consistency loop in DMFT. There are essentially two options: to

close it directly in complex time or to perform an analytic continuation in every iteration. Conceptually, the first option offers many advantages and must be explored rigorously. Nonetheless, further improvements of the post-processing of the tilted contour are equally important, especially as complex time evolution is not restricted to impurity models.

Our advances enabled us to explore intricate physical phenomena at very low temperatures. Combined with our efficient MT3N multi-orbital impurity solver these improvements will allow us to accurately treat full 5-orbital systems, or even 7-orbital systems in the near future. Additionally, we want to stress the relative simplicity and stability of extending the impurity Hamiltonian in TN-based impurity solvers. This allows us to incorporate additional interactions effects present in materials. Our recent advances in the simulation of electron-phonon systems [3, 7] are transferable to our impurity solver. Incorporating them into DFT+DMFT studies could mark an important step towards a better understanding of the role of phonons in real materials.

Pushing the boundaries on what is numerically feasible inevitably brings us closer to regions where the single-site approximation of DMFT eventually breaks down. Therefore, it is important to extend efforts to cluster-DMFT methods, also in the context of material research. While cellular DMFT methods sparked more interest recently [4], we consider the possible extension of the DFT+DMFT approach to DFT+DCA (dynamical cluster approximation) to be similarly intriguing, specifically considering that one will be severely limited in the cluster size initially.

Appendix A

Post-processing for Complex Time Contours

The following results are based on the author's publication [6].

The ideal choice of the complex contour ultimately relies on the model and its specific high energy quasi-particle peak structure because the loss of high-energy information is greater if the contour is further away from the real-time axis during early times. A second crucial aspect is the growth of entanglement entropies on the complex contours, as it determines the computational cost. This is illustrated on figure A.1. Not surprisingly, the real-time contour has the strongest overall growth of entanglement, which does not saturate, explaining the vastly larger numerical resources required. The tilted contour (here $\alpha = 0.1$) starts with a similar growth of entanglement at early times when it is close to the real axis but actually reveals even a slight decrease before saturation at some relatively small value. The kink contour without post-processing (kink $\alpha = 0.1$) has a very strong initial growth of entanglement, even stronger than for real-time evolution because it does not exploit time-splitting in the real-time direction. Ultimately, though, entanglement strongly resembles that of the tilted contour, reminiscent of the observed saturation of entanglement growth for all complex-time contours. Finally, the parallel contour at constant τ has the slowest early entanglement growth as it aggressively suppresses high-energy states but settles on a somewhat larger saturation value as the contour stays closer to the real axis. (Note that the $\alpha = 0.1$ and $\tau = 0.563$ data do not allow for immediate comparison; they are rather indicative of typical behavior for the respective contours.) The saturation of entanglement observed for all complex contours explains the large accessible times.




Figure A.1: Comparison of the entanglement entropy on different complex-time contours. We show the sum of the entanglement entropies for all system cuts as the system is heterogeneous. Kink $\alpha = 0.1$ is the post-processing free contour with vertical contour along the negative imaginary axis and a tilted contour in a positive imaginary direction until $t = 30D^{-1}$ and $\tau = 0.563 \text{ eV}^{-1}$ a parallel contour with constant imaginary part $\tau = 0.563 \text{ eV}^{-1}$. Times of contours exploiting real-time splitting were multiplied by a factor of 2 to represent the same final time.

A.1 Extrapolation Method

We assume that we know n Green's functions $G(t + i\tau_k)$ for n different τ_k . As seen previously, they can be generated at low numerical cost from a single contour. We approximate the behavior by a power series incorporating terms up to order $n - 1$,

$$G(t + i\tau) = G(t) + \sum_{m=1}^{n-1} \tau^m h_m(t). \quad (\text{A.1})$$

For n values τ_k , we have n equations with n unknown variables, the Green's function $G(t)$ we are interested in, and $n - 1$ coefficients $h_m(t)$ which we will not require explicitly. This linear equation system can be written as $\mathbf{g} = \mathbf{M} \cdot \mathbf{h}$, where the vector components are $g_k = G(t + i\tau_k)$, $h_0 = G(t)$ and $h_k = h_k(t)$ for $k > 0$, and

$$M_{kl} = \tau_k^{l-1}. \quad (\text{A.2})$$

An inversion of \mathbf{M} yields

$$G(t) = \sum_{k=1}^n (M_{1k})^{-1} G(t + i\tau_k). \quad (\text{A.3})$$

For increasing n , the inversion of \mathbf{M} becomes more difficult. To stabilize it, it is useful to re-scale both $h_k(t)$ and \mathbf{M} such that $G(t)$ remains unaffected. With τ_{\max} the largest τ_k , we re-scale as

$$M_{kl} = \left(\frac{\tau_k}{\tau_{\max}} \right)^{l-1}. \quad (\text{A.4})$$

If we want to solve at order n , but know $G(t + i\tau_k)$ for more than n different τ_k , we found a substantial increase of accuracy by averaging over all different n th-order extrapolations with different choices of n values of τ_k . In practice, we used $n = 6$ for 13 different $G(t + i\tau_k)$. In the calculations shown in the main text (figure 7.5 and figure 7.4), they were chosen equidistantly ($\delta\tau = 0.075$) centered on $\tau = 1$. (Note that the extrapolation method can also be applied to the linear contours of section 7.2, see figure 7.12.)

A.2 Stabilizing the Kink Method

The kink method appears rather charming due to the absence of any analytic continuation. However, the observed breakdown of accuracy strongly limits its efficacy. We believe that this breakdown is the result of the numerically challenging time evolution towards negative complex time. This evolution increases the weight of high energy contributions in the state and slowly moves towards the bulk. Thus, we expect that the area law of quantum states no longer holds, and we find a dependence on the system size in the breakdown behavior. Indeed if we compute the overlap of a quantum state given as

$$\langle \psi_0 | c^\dagger(0)c(0) | \psi_0 \rangle = [c(i\tau) | \psi_0 \rangle]^\dagger c(i\tau) | \psi_0 \rangle \quad (\text{A.5})$$

with τ in $[0, \tau_{\max}]$ we expect it to be constant for all τ . However, the ill-conditioned nature of the time evolution towards positive complex time leads to a violation of this equation, see figure A.2. We can see a clear dependence in both the error in the Green's function as in the slope of the increasing entanglement entropy summed over all bonds. This indicates that the breakdown of this method is directly linked to the breakdown of the tensor network approximation, i.e. the inability to efficiently represent bulk states.




Figure A.2: Breakdown of the overlap defined in equation (A.5) for different system sizes. We find a dependence on the system size in the error of the overlap (left) and in the slope of the increasing entanglement entropy (right). Note the stark increase of the entanglement entropy in the time evolution towards negative τ (colored) compared to the growth in the evolution towards positive τ for $L = 100$ (gray).

Appendix B

LiV₂O₄: Comparison of Multi-orbital models

While a full 5-orbital treatment of this material is unnecessary to probe its intricate low-energy physics, it is nonetheless a testament to the capabilities of our impurity solver. The e_g orbitals are far removed from the Fermi edge, resulting in a minuscule electronic occupation in our DMFT study of $n_{e_g} = 0.01$, for zero temperature simulations at $\beta_{\text{eff}} = 1000$. Despite this seemingly small orbital occupation, the inclusion of the e_g orbitals increases the occupation of the a_{1g} orbital even further, from $n_{a_{1g}} = 0.893$ to $n_{a_{1g}} = 0.918$. Note that both these values are for zero-temperature simulations.

Note that neither *even* nor *odd* bath sizes result in an S=0 ground state for our 5-orbital simulations; we hence cannot compare this compound's physically relevant low-energy behavior. While this is arguably simply a bias in our discretization procedure which tends to overestimate the local magnetic moment, it is also a testament to the strong competition of correlation effects at such low temperatures.

The 5-orbital model shows a slightly increased scattering rate compared to the 3-orbital simulation, which we consider to be a mere consequence of the larger orbital occupation, as shown in figure B.1. The e_g^π orbitals are remarkably similar in both descriptions, while the e_g orbitals show only weak renormalization effects. The latter is to be expected given how far removed the e_g orbitals are from the Fermi edge.




Figure B.1: Comparison of self-energies from 3-orbital and 5-orbital simulations of LiV_2O_4 . The self-energy of the a_{1g} orbital (left) reveals an increased scattering rate which we contribute to the increased orbital occupation. The e_g^π orbitals (right) are remarkably similar compared to the 3-orbital simulation. We also show the self-energy of the e_g orbital (right), displaying only small renormalization effects congruent with its electronic occupation.

List of Figures

2.1	Decomposition of the coefficient tensor from the left	9
2.2	Representation of a mixed canonical MPS	10
2.3	Representation of a mixed canonical MPS	11
2.4	Maximum bond dimension profile of a MPS	12
2.5	Graphical depiction of leg directions on MPS	13
2.6	Graphical depiction of a tensor map	13
2.7	Graphical depiction of the quantum number flow in MPS	14
2.8	Effective single-site Hamiltonian	15
2.9	DMRG single site update	16
2.10	Expansion Tensor	17
2.11	Comparison of different TDVP algorithms based on LiV_2O_4	21
3.1	Mapping from the lattice to the impurity model	28
3.2	Splitting of d orbital levels due to lattice symmetries	32
4.1	Depiction of finite temperature and $T = 0$ time-ordered Green's function	41
4.2	Zero temperature time-ordered Green's function	42
4.3	Zero temperature impurity Green's function on the Matsubara axis	44
4.4	Zero temperature self-energy on the Matsubara axis	44
4.5	Reconstructed time-ordered Green's function from linear prediction data	45
4.6	Zero temperature comparison of self energy with linear prediction from different τ	49
4.7	Comparison of old and new IT recipe	50
4.8	Comparison of dynamic quantities with linear and exponential time steps	51
4.9	Sketch of line-fit and inflection point optimization for α in AC	54
4.10	Comparison between MaxEnt and MinKL for the DN Hamiltonian	57
4.11	Comparison between MaxEnt and MinKL for LiV_2O_4	58
4.12	α dependency of MinKL analytic continuation	59
4.13	Analytic Continuation of $\Sigma(i\omega_n)$ for LiV_2O_4	61
5.1	Convergence of discretization procedure with bath sites	65
5.2	Cost function comparisons for SIAM and LiV_2O_4	67
5.3	Revival times for equally spaced intervals	68

5.4	Revival times for different discretization techniques	69
5.5	Fourier-domain discretization procedure for the SIAM	70
5.6	Fourier-domain discretization procedure for LiV_2O_4	71
5.7	Cost functions of Fourier domain discretization procedure for LiV_2O_4	72
5.8	Convergence speed of impurity Green's function in the SIAM model for different discretization procedures	74
5.9	Convergence speed of self-energy in the SIAM model for different discretization procedures	74
5.10	The hopping structure of an impurity model in star-geometry.	76
5.11	The hopping structure of an impurity model in chain-geometry.	77
5.12	Comparison of ground state occupations in star- and chain-geometry representation	78
5.13	One particle reduced density matrix comparison between star- and chain-geometry	79
5.14	The hopping structure of an impurity model in double-chain mapping.	81
5.15	One particle reduced density matrix and occupation structure of double chain-geometry	82
5.16	Entanglement growth over time of the SIAM model in different representations	83
5.17	Mapping keeping spin degrees of freedom close to each other	84
5.18	Mapping separating spin degrees of freedom	84
5.19	Runtime comparison between the iterative ground state search technique and a full ground state search based upon a particle number guess	86
5.20	Overlap between ground states and different initial states in the SIAM	87
5.21	Splitting of real time evolution in the SIAM	90
6.1	Depiction of the scaling for different updates in the FTNS	92
6.2	Depiction of a three-legged tree tensor network	93
6.3	Depiction of a fork three-legged tree tensor network	94
6.4	Depiction of a minimal three-legged tree tensor network	94
6.5	Scaling of TNS solvers with number of orbitals	97
6.6	Comparison MT3N vs FTNS SrMoO_3	99
7.1	Light cone of an electron during complex time evolution	103
7.2	Depiction of parallel contour	104
7.3	Bond dimension profile during complex time evolution	105
7.4	Comparison for different orders of the superposed inversion method	107
7.5	Comparison of spectral functions obtained on a parallel contour	108
7.6	Comparison of spectral functions and self-energies obtained on a parallel contour for various U	109
7.7	Symmetrically split complex-time contour	110
7.8	Comparison of spectral functions obtained on tilted contours	111
7.9	Energy and entanglement entropy during time-evolution	111
7.10	Contour for complex time evolution with no post-processing	113

7.11	Spectral functions on the kink contour	114
7.12	Spectral function $\mathcal{A}(\omega)$ for a three-orbital model with Hubbard–Kanamori interaction	115
7.13	Self-energy of the Hubbard–Kanamori Hamiltonian	116
7.14	Self-energy of the Dworin–Narath Hamiltonian	118
8.1	DFT band structure of LiV_2O_4	124
8.2	Comparison of self-energies for $\beta = 1000$ and $\beta_{\text{eff}} = 1000$ results in LiV_2O_4	127
8.3	Infinite temperature MPS in the purification framework	128
8.4	Formation of a gap at $\tau = \beta/2$ in the time ordered Green’s function	130
8.5	Gap in finite temperature simulations	131
8.6	Mitigation of oscillations in finite temperature self-energy	131
8.7	Impact of the <i>odd/even</i> effect on LiV_2O_4 orbitals	133
8.8	Self energy for different bath sizes at various total spins	134
8.9	Emergence of heavy Fermi liquid behaviour in LiV_2O_4	135
8.10	Self-energy for various Hund’s coupling strengths	136
8.11	Spin freezing regime of the Hubbard-Kanamori Hamiltonian	138
8.12	Local spin-spin susceptibility in LiV_2O_4	139
A.1	Comparison of the entanglement entropy growth of different complex contours	146
A.2	Breakdown of the kink method	148
B.1	Comparison of 3-orbital and 5-orbital simulations of LiV_2O_4	150

List of Tables

3.1	Comparison of impurity solvers	34
5.1	Comparison of mutual information and overlaps between a Fermi-state and the actual ground state for different mappings of the SIAM model	81
6.1	Comparison of runtimes for different solvers during DMRG	96
6.2	Simulation parameters for IT solvers	98
6.3	Comparison of errors for different solvers	98
6.4	Runtime comparison for different TDVP modes for the FTSP and MT3N structures	100

Bibliography

- [1] J. Karp, M. Bramberger, M. Grundner, U. Schollwöck, A. J. Millis, and M. Zingl, Phys. Rev. Lett. **125**, 166401 (2020).
- [2] M. Bramberger, J. Mravlje, M. Grundner, U. Schollwöck, and M. Zingl, Phys. Rev. B **103**, 165133 (2021).
- [3] S. Mardazad, Y. Xu, X. Yang, M. Grundner, U. Schollwöck, H. Ma, and S. Paeckel, The Journal of Chemical Physics **155**, 194101 (2021).
- [4] M. Bramberger, B. Bacq-Labreuil, M. Grundner, S. Biermann, U. Schollwöck, S. Paeckel, and B. Lenz, SciPost Phys. **14**, 010 (2023).
- [5] M. Moroder, M. Grundner, F. Damanet, U. Schollwöck, S. Mardazad, S. Flannigan, T. Köhler, and S. Paeckel, Physical Review B **107**, 10.1103/physrevb.107.214310 (2023).
- [6] M. Grundner, P. Westhoff, F. B. Kugler, O. Parcollet, and U. Schollwöck, Phys. Rev. B **109**, 155124 (2024).
- [7] M. Grundner, T. Blatz, J. Sous, U. Schollwöck, and S. Paeckel, Cooper-paired bipolaronic superconductors (2023), arXiv:2308.13427 .
- [8] M. Grundner, F. B. Kugler, O. Parcollet, U. Schollwöck, A. Georges, and A. Hampel, (2024), arXiv:2409.17268 [cond-mat.str-el] .
- [9] F. Bloch, Zeitschrift für Physik **52**, 555 (1929).
- [10] J. Lilienfeld, US Patent (1930).
- [11] S. Yuan, X. Duan, J. Liu, Y. Ye, F. Lv, T. Liu, Q. Wang, and X. Zhang, Energy Storage Materials **42**, 317 (2021).
- [12] D. N. Basov, R. D. Averitt, and D. Hsieh, Nature Materials **16**, 1077 (2017).
- [13] A. Hampel, *Interplay between structural, electronic, and magnetic properties in rare-earth nickelates*, Ph.D. thesis, ETH Zurich (2019).
- [14] F. M. Granozio, G. Koster, and G. Rijnders, MRS Bulletin **38**, 1017 (2013).

- [15] C. Ahn, A. Cavalleri, A. Georges, S. Ismail-Beigi, A. J. Millis, and J.-M. Triscone, *Nature Materials* **20**, 1462 (2021).
- [16] L. Ye, S. Fang, M. Kang, J. Kaufmann, Y. Lee, C. John, P. M. Neves, S. Y. F. Zhao, J. Denlinger, C. Jozwiak, A. Bostwick, E. Rotenberg, E. Kaxiras, D. C. Bell, O. Janson, R. Comin, and J. G. Checkelsky, *Nature Physics* 10.1038/s41567-023-02360-5 (2024).
- [17] W. S. Choi, C. M. Rouleau, S. S. A. Seo, Z. Luo, H. Zhou, T. T. Fister, J. A. Eastman, P. H. Fuoss, D. D. Fong, J. Z. Tischler, G. Eres, M. F. Chisholm, and H. N. Lee, *Advanced Materials* **24**, 6423 (2012), <https://onlinelibrary.wiley.com/doi/pdf/10.1002/adma.201202691> .
- [18] Q. Lei, M. Golalikhani, B. A. Davidson, G. Liu, D. G. Schlom, Q. Qiao, Y. Zhu, R. U. Chandrasena, W. Yang, A. X. Gray, E. Arenholz, A. K. Farrar, D. A. Tenne, M. Hu, J. Guo, R. K. Singh, and X. Xi, *npj Quantum Materials* **2**, 10 (2017).
- [19] J. Hubbard, *Proc. R. Soc. Lond. A* , 238 (1963).
- [20] P. Hohenberg and W. Kohn, *Phys. Rev.* **136**, B864 (1964).
- [21] W. Kohn and L. J. Sham, *Phys. Rev.* **140**, A1133 (1965).
- [22] W. Kohn, *Rev. Mod. Phys.* **71**, 1253 (1999).
- [23] M. Imada, A. Fujimori, and Y. Tokura, *Rev. Mod. Phys.* **70**, 1039 (1998).
- [24] V. I. Anisimov, J. Zaanen, and O. K. Andersen, *Phys. Rev. B* **44**, 943 (1991).
- [25] B. Himmetoglu, A. Floris, S. de Gironcoli, and M. Cococcioni, *International Journal of Quantum Chemistry* **114**, 14 (2014), <https://onlinelibrary.wiley.com/doi/pdf/10.1002/qua.24521> .
- [26] J. M. Tomczak, P. Liu, A. Toschi, G. Kresse, and K. Held, *The European Physical Journal Special Topics* **226**, 2565 (2017).
- [27] A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg, *Rev. Mod. Phys.* **68**, 13 (1996).
- [28] A. Georges, L. d. Medici, and J. Mravlje, *Annual Review of Condensed Matter Physics* **4**, 137 (2013), <https://doi.org/10.1146/annurev-conmatphys-020911-125045> .
- [29] G. Kotliar, S. Y. Savrasov, K. Haule, V. S. Oudovenko, O. Parcollet, and C. A. Marianetti, *Rev. Mod. Phys.* **78**, 865 (2006).
- [30] K. Held, *Advances in Physics* **56**, 829 (2007), <https://doi.org/10.1080/00018730701619647> .

- [31] C. Martins, M. Aichhorn, L. Vaugier, and S. Biermann, Phys. Rev. Lett. **107**, 266404 (2011).
- [32] I. Leonov, V. I. Anisimov, and D. Vollhardt, Phys. Rev. Lett. **112**, 146401 (2014).
- [33] N.-O. Linden, M. Zingl, C. Hubig, O. Parcollet, and U. Schollwöck, Phys. Rev. B **101**, 041101 (2020).
- [34] M. Bramberger, *Effects of electronic correlations in BaOsO₃ and tetragonal CuO*, Ph.D. thesis, Ludwig-Maximilians-Universität München (2022).
- [35] F. A. Wolf, A. Go, I. P. McCulloch, A. J. Millis, and U. Schollwöck, Phys. Rev. X **5**, 041032 (2015).
- [36] F. A. Wolf, I. P. McCulloch, and U. Schollwöck, Phys. Rev. B **90**, 235131 (2014).
- [37] F. A. Wolf, I. P. McCulloch, O. Parcollet, and U. Schollwöck, Phys. Rev. B **90**, 115124 (2014).
- [38] M. Ganahl, M. Aichhorn, H. G. Evertz, P. Thunström, K. Held, and F. Verstraete, Phys. Rev. B **92**, 155132 (2015).
- [39] D. Bauernfeind, M. Zingl, R. Triebl, M. Aichhorn, and H. G. Evertz, Phys. Rev. X **7**, 031013 (2017).
- [40] M. Caffarel and W. Krauth, Phys. Rev. Lett. **72**, 1545 (1994).
- [41] Y. Lu, M. Höppner, O. Gunnarsson, and M. W. Haverkort, Phys. Rev. B **90**, 085102 (2014).
- [42] K. G. Wilson, Rev. Mod. Phys. **47**, 773 (1975).
- [43] R. Bulla, T. Costi, and T. Pruschke, Reviews of Modern Physics, vol. 80, Issue 2, pp. 395-450 (2008).
- [44] F. B. Kugler, M. Zingl, H. U. R. Strand, S.-S. B. Lee, J. von Delft, and A. Georges, Phys. Rev. Lett. **124**, 016401 (2020).
- [45] E. Gull, A. J. Millis, A. I. Lichtenstein, A. N. Rubtsov, M. Troyer, and P. Werner, Rev. Mod. Phys. **83**, 349 (2011).
- [46] S. Kondo, D. C. Johnston, C. A. Swenson, F. Borsa, A. V. Mahajan, L. L. Miller, T. Gu, A. I. Goldman, M. B. Maple, D. A. Gajewski, E. J. Freeman, N. R. Dilley, R. P. Dickey, J. Merrin, K. Kojima, G. M. Luke, Y. J. Uemura, O. Chmaissem, and J. D. Jorgensen, Phys. Rev. Lett. **78**, 3729 (1997).
- [47] I. V. Oseledets, SIAM Journal on Scientific Computing **33**, 2295 (2011), <https://doi.org/10.1137/090752286> .

- [48] I. Oseledets and E. Tyrtysnikov, *Linear Algebra and its Applications* **432**, 70 (2010).
- [49] D. V. Savostyanov, *Linear Algebra and its Applications* **458**, 217 (2014).
- [50] S. Dolgov and D. Savostyanov, *Computer Physics Communications* **246**, 106869 (2020).
- [51] L. I. Vysotsky, A. V. Smirnov, and E. E. Tyrtysnikov, *Lobachevskii Journal of Mathematics* **42**, 1608 (2021).
- [52] Y. Núñez Fernández, M. Jeannin, P. T. Dumitrescu, T. Kloss, J. Kaye, O. Parcollet, and X. Waintal, *Phys. Rev. X* **12**, 041018 (2022).
- [53] H. Shinaoka, M. Wallerberger, Y. Murakami, K. Nogaki, R. Sakurai, P. Werner, and A. Kauch, *Phys. Rev. X* **13**, 021015 (2023).
- [54] S. V. Dolgov, B. N. Khoromskij, and I. V. Oseledets, *SIAM Journal on Scientific Computing* **34**, A3016 (2012), <https://doi.org/10.1137/120864210> .
- [55] M. K. Ritter, Y. N. Fernández, M. Wallerberger, J. von Delft, H. Shinaoka, and X. Waintal, *Quantics tensor cross interpolation for high-resolution, parsimonious representations of multivariate functions in physics and beyond* (2023), arXiv:2303.11819 .
- [56] V. Kazeev, M. Khammash, M. Nip, and C. Schwab, *PLOS Computational Biology* **10**, 1 (2014).
- [57] N. Gourianov, M. Lubasch, S. Dolgov, Q. Y. van den Berg, H. Babae, P. Givi, M. Kiffner, and D. Jaksch, *Nature Computational Science* **2**, 30 (2022).
- [58] M. Kiffner and D. Jaksch, *Tensor network reduced order models for wall-bounded flows* (2023), arXiv:2303.03010 .
- [59] K. Sozykin, A. Chertkov, R. Schutski, A.-H. Phan, A. Cichocki, and I. Oseledets, *Ttopt: A maximum volume quantized tensor train-based optimization and its application to reinforcement learning* (2022), arXiv:2205.00293 .
- [60] A. Tomut, S. S. Jahromi, S. Singh, F. Ishtiaq, C. Muñoz, P. S. Bajaj, A. Elborady, G. del Bimbo, M. Alizadeh, D. Montero, P. Martin-Ramiro, M. Ibrahim, O. T. Alaoui, J. Malcolm, S. Mugel, and R. Orus, *Compactifai: Extreme compression of large language models using quantum-inspired tensor networks* (2024), arXiv:2401.14109 .
- [61] H. Bethe, *Zeitschrift für Physik* **71**, 205 (1931).
- [62] N. Andrei, *Phys. Rev. Lett.* **45**, 379 (1980).
- [63] S. R. White, *Phys. Rev. Lett.* **69**, 2863 (1992).

- [64] S. R. White, Phys. Rev. B **48**, 10345 (1993).
- [65] S. Rommer and S. Östlund, Phys. Rev. B **55**, 2164 (1997).
- [66] E. Jeckelmann, Phys. Rev. B **66**, 045114 (2002).
- [67] U. Schollwöck, Rev. Mod. Phys. **77**, 259 (2005).
- [68] S. Mardazad, *Simulating real molecules with tensor network techniques*, Ph.D. thesis, Ludwig-Maximilians-Universität München (2022).
- [69] K. Gunst, F. Verstraete, S. Wouters, Ö. Legeza, and D. Van Neck, Journal of Chemical Theory and Computation **14**, 2026 (2018).
- [70] F. Verstraete, M. M. Wolf, D. Perez-Garcia, and J. I. Cirac, Phys. Rev. Lett. **96**, 220601 (2006).
- [71] J. Tindall, M. Fishman, E. M. Stoudenmire, and D. Sels, PRX Quantum **5**, 10.1103/prxquantum.5.010308 (2024).
- [72] E. Stoudenmire and S. R. White, Annual Review of Condensed Matter Physics **3**, 111–128 (2012).
- [73] U. Schollwöck, Annals of Physics **326**, 96 (2011), january 2011 Special Issue.
- [74] M. B. Hastings, Journal of Statistical Mechanics: Theory and Experiment **2007**, P08024 (2007).
- [75] J. Eisert, M. Cramer, and M. B. Plenio, Rev. Mod. Phys. **82**, 277 (2010).
- [76] T. J. Osborne, Phys. Rev. A **75**, 032321 (2007).
- [77] C. Hubig, F. Lachenmaier, N.-O. Linden, T. Reinhard, L. Stenzel, A. Swoboda, M. Grundner, and S. Mardazad, The SYTEN toolkit.
- [78] C. Hubig, *Symmetry-protected tensor networks*, Ph.D. thesis, Ludwig-Maximilians-Universität München (2017).
- [79] I. P. McCulloch, Infinite size density matrix renormalization group, revisited (2008), arXiv:0804.2509 .
- [80] S. R. White, Phys. Rev. B **72**, 180403 (2005).
- [81] Y. Núñez Fernández and G. Torroba, Phys. Rev. B **101**, 085135 (2020).
- [82] T. Köhler, J. Stolpp, and S. Paeckel, SciPost Phys. **10**, 058 (2021).
- [83] C. Hubig, I. P. McCulloch, U. Schollwöck, and F. A. Wolf, Physical Review B **91**, 10.1103/physrevb.91.155115 (2015).

- [84] S. V. Dolgov and D. V. Savostyanov, *SIAM Journal on Scientific Computing* **36**, A2248 (2014), <https://doi.org/10.1137/140953289> .
- [85] S. V. Dolgov and D. V. Savostyanov, in *Numerical Mathematics and Advanced Applications - ENUMATH 2013*, edited by A. Abdulle, S. Deparis, D. Kressner, F. Nobile, and M. Picasso (Springer International Publishing, Cham, 2015) pp. 335–343.
- [86] P. W. Langhoff, S. T. Epstein, and M. Karplus, *Rev. Mod. Phys.* **44**, 602 (1972).
- [87] A. McLachlan, *Molecular Physics* **8**, 39 (1964), <https://doi.org/10.1080/00268976400100041> .
- [88] P. A. M. Dirac, *Mathematical Proceedings of the Cambridge Philosophical Society* **26**, 376–385 (1930).
- [89] N. Mott, *The Mathematical Gazette* **17**, 135–136 (1933).
- [90] J. Haegeman, C. Lubich, I. Oseledets, B. Vandereycken, and F. Verstraete, *Phys. Rev. B* **94**, 165116 (2016).
- [91] C. Lubich, *Math. Comput.* **74**, 765 (2004).
- [92] S. Paeckel, T. Köhler, A. Swoboda, S. R. Manmana, U. Schollwöck, and C. Hubig, *Annals of Physics* **411**, 167998 (2019).
- [93] H.-D. Meyer, U. Manthe, and L. Cederbaum, *Chemical Physics Letters* **165**, 73 (1990).
- [94] U. Manthe, H. Meyer, and L. S. Cederbaum, *The Journal of Chemical Physics* **97**, 3199 (1992), https://pubs.aip.org/aip/jcp/article-pdf/97/5/3199/11323495/3199_1_online.pdf .
- [95] M. Beck, A. Jäckle, G. Worth, and H.-D. Meyer, *Physics Reports* **324**, 1 (2000).
- [96] M. Hochbruck and C. Lubich, *SIAM Journal on Numerical Analysis* **34**, 1911 (1997).
- [97] M. Hochbruck and C. Lubich, *BIT* **39**, 620–645 (1999).
- [98] S. Paeckel, T. Köhler, S. R. Manmana, and B. Lenz, *Matrix-product-state-based band-lanczos solver for quantum cluster approaches* (2023), arXiv:2310.10799 .
- [99] C. Hubig, J. Haegeman, and U. Schollwöck, *Physical Review B* **97**, 10.1103/physrevb.97.045125 (2018).
- [100] C. Bertrand, D. Bauernfeind, P. T. Dumitrescu, M. Maček, X. Waintal, and O. Parcollet, *Phys. Rev. B* **103**, 155104 (2021).
- [101] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, *Phys. Rev. Lett.* **100**, 030504 (2008).

- [102] P. Calabrese and J. Cardy, *Journal of Statistical Mechanics: Theory and Experiment* **2004**, P06002 (2004).
- [103] P. Calabrese and J. Cardy, *Journal of Statistical Mechanics: Theory and Experiment* **2005**, P04010 (2005).
- [104] P. Calabrese and J. Cardy, *Journal of Statistical Mechanics: Theory and Experiment* **2007**, P10004 (2007).
- [105] J. Eisert and T. J. Osborne, *Phys. Rev. Lett.* **97**, 150404 (2006).
- [106] T. J. Osborne, *Phys. Rev. Lett.* **97**, 157202 (2006).
- [107] M. Yang and S. R. White, *Physical Review B* **102**, 10.1103/physrevb.102.094315 (2020).
- [108] T. Blatz, *Subspace Expansion Methods for Matrix Product States*, Master's thesis, Ludwig-Maximilians-Universität München (2022).
- [109] M. Yang and S. R. White, Time dependent variational principle with ancillary krylov subspace (2020).
- [110] G. Kotliar, S. Y. Savrasov, G. Pálsson, and G. Biroli, *Phys. Rev. Lett.* **87**, 186401 (2001).
- [111] S. Biermann, F. Aryasetiawan, and A. Georges, *Phys. Rev. Lett.* **90**, 086402 (2003).
- [112] N. Lin, C. A. Marianetti, A. J. Millis, and D. R. Reichman, *Phys. Rev. Lett.* **106**, 096402 (2011).
- [113] T. Zhu, Z.-H. Cui, and G. K.-L. Chan, *Journal of Chemical Theory and Computation* **16**, 141 (2020), pMID: 31815457, <https://doi.org/10.1021/acs.jctc.9b00934> .
- [114] W. Metzner and D. Vollhardt, *Phys. Rev. Lett.* **62**, 324 (1989).
- [115] A. Georges and G. Kotliar, *Phys. Rev. B* **45**, 6479 (1992).
- [116] A. Georges and W. Krauth, *Phys. Rev. Lett.* **69**, 1240 (1992).
- [117] Y. Kuramoto and T. Watanabe, *Physica B+C* **148**, 80 (1987).
- [118] I. H. Inoue, I. Hase, Y. Aiura, A. Fujimori, Y. Haruyama, T. Maruyama, and Y. Nishihara, *Phys. Rev. Lett.* **74**, 2539 (1995).
- [119] M. H. Hettler, A. N. Tahvildar-Zadeh, M. Jarrell, T. Pruschke, and H. R. Krishnamurthy, *Phys. Rev. B* **58**, R7475 (1998).
- [120] M. H. Hettler, M. Mukherjee, M. Jarrell, and H. R. Krishnamurthy, *Phys. Rev. B* **61**, 12739 (2000).

- [121] D. Sénéchal, D. Perez, and M. Pioro-Ladrière, Phys. Rev. Lett. **84**, 522 (2000).
- [122] J. Steinbauer, L. de' Medici, and S. Biermann, Phys. Rev. B **100**, 085104 (2019).
- [123] R. Žitko, Physical Review B **80**, 10.1103/physrevb.80.125125 (2009).
- [124] T. Pruschke, R. Bulla, and M. Jarrell, Phys. Rev. B **61**, 12799 (2000).
- [125] T. Pruschke, R. Bulla, and M. Jarrell, Physica B: Condensed Matter **281-282**, 47 (2000).
- [126] P. W. Anderson, Phys. Rev. **124**, 41 (1961).
- [127] F. Aryasetiawan, M. Imada, A. Georges, G. Kotliar, S. Biermann, and A. I. Lichtenstein, Phys. Rev. B **70**, 195104 (2004).
- [128] E. Şaşoğlu, C. Friedrich, and S. Blügel, Phys. Rev. B **83**, 121101 (2011).
- [129] S. Sugano, Y. Tanabe, and H. Kamimura, eds., *Multiplets of Transition-Metal Ions in Crystals*, Vol. 33 (New York: Academic Press, 1970).
- [130] L. Dworin and A. Narath, Phys. Rev. Lett. **25**, 1287 (1970).
- [131] P. Werner and A. J. Millis, Phys. Rev. B **74**, 155107 (2006).
- [132] P. Werner, A. Comanac, L. de' Medici, M. Troyer, and A. J. Millis, Phys. Rev. Lett. **97**, 076405 (2006).
- [133] M. Kim, J. Mravlje, M. Ferrero, O. Parcollet, and A. Georges, Physical review letters **120 12**, 126401 (2017).
- [134] D. Sutter, C. G. Fatuzzo, S. Moser, M. Kim, R. Fittipaldi, A. Vecchione, V. Granata, Y. Sassa, F. Cossalter, G. Gatti, M. Grioni, H. M. Rønnow, N. C. Plumb, C. E. Matt, M. Shi, M. Hoesch, T. K. Kim, T.-R. Chang, H.-T. Jeng, C. Jozwiak, A. Bostwick, E. Rotenberg, A. Georges, T. Neupert, and J. Chang, Nature Communications **8**, 15176 (2017).
- [135] S. Beck, A. Hampel, O. Parcollet, C. Ederer, and A. Georges, Journal of Physics: Condensed Matter **34**, 235601 (2022).
- [136] E. Y. Loh, J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino, and R. L. Sugar, Phys. Rev. B **41**, 9301 (1990).
- [137] M. Troyer and U.-J. Wiese, Phys. Rev. Lett. **94**, 170201 (2005).
- [138] A. J. Kim, P. Werner, and R. Valentí, Phys. Rev. B **101**, 045108 (2020).
- [139] J. Vučičević, J. Kokalj, R. Žitko, N. Wentzell, D. Tanasković, and J. Mravlje, Phys. Rev. Lett. **123**, 036601 (2019).

- [140] J. Karp, A. Hampel, M. Zingl, A. S. Botana, H. Park, M. R. Norman, and A. J. Millis, *Phys. Rev. B* **102**, 245130 (2020).
- [141] A. Horvat, R. Žitko, and J. Mravlje, *Phys. Rev. B* **94**, 165140 (2016).
- [142] A. Horvat, R. Žitko, and J. Mravlje, *Phys. Rev. B* **96**, 085122 (2017).
- [143] C. Raas, G. S. Uhrig, and F. B. Anders, *Phys. Rev. B* **69**, 041102 (2004).
- [144] C. Raas and G. S. Uhrig, *The European Physical Journal B - Condensed Matter and Complex Systems* **45**, 293 (2005).
- [145] M. Ganahl, P. Thunström, F. Verstraete, K. Held, and H. G. Evertz, *Phys. Rev. B* **90**, 045144 (2014).
- [146] T. Barthel, U. Schollwöck, and S. R. White, *Physical Review B* **79**, 10.1103/physrevb.79.245101 (2009).
- [147] Y. Tian and S. R. White, Matrix product state recursion methods for strongly correlated quantum systems (2020), arXiv:2010.00213 .
- [148] L. Lin and S. R. White, A novel method of function extrapolation inspired by techniques in low-entangled many-body physics (2023), arXiv:2308.09001 .
- [149] B. Y. Hu, *American Journal of Physics* **57**, 821 (1989), https://pubs.aip.org/aapt/ajp/article-pdf/57/9/821/11438014/821_1_online.pdf .
- [150] G. J. Krabberger and M. Zingl, *Triqs/maxent*.
- [151] E. T. Jaynes, *Phys. Rev.* **106**, 620 (1957).
- [152] J. Skilling, *Classic maximum entropy* (Springer, 1989) pp. 45–52.
- [153] R. N. Silver, D. S. Sivia, and J. E. Gubernatis, *Phys. Rev. B* **41**, 2380 (1990).
- [154] K. S. D. Beach, R. J. Gooding, and F. Marsiglio, *Phys. Rev. B* **61**, 5147 (2000).
- [155] J. Fei, C.-N. Yeh, and E. Gull, *Phys. Rev. Lett.* **126**, 056402 (2021).
- [156] D. Bergeron and A.-M. Tremblay, *Physical Review E* **94**, 023303 (2016).
- [157] G. J. Krabberger, R. Triebl, M. Zingl, and M. Aichhorn, *Phys. Rev. B* **96**, 155128 (2017).
- [158] J. Mravlje and A. Georges, *Phys. Rev. Lett.* **117**, 036401 (2016).
- [159] R. K. Bryan (Kluwer Academic Publishers, 1990) pp. 221—232.
- [160] M. Grundner, P. Westhoff, S. Paeckel, and U. Schollwöck (in prep.).

- [161] S. Kullback and R. A. Leibler, *The Annals of Mathematical Statistics* **22**, 79 (1951).
- [162] J. E. Gubernatis, M. Jarrell, R. N. Silver, and D. S. Sivia, *Phys. Rev. B* **44**, 6011 (1991).
- [163] A. Shimoyamada, S. Tsuda, K. Ishizaka, T. Kiss, T. Shimojima, T. Togashi, S. Watanabe, C. Q. Zhang, C. T. Chen, Y. Matsushita, H. Ueda, Y. Ueda, and S. Shin, *Phys. Rev. Lett.* **96**, 026403 (2006).
- [164] J. Kaye, K. Chen, and H. U. Strand, *Computer Physics Communications* **280**, 108458 (2022).
- [165] J. Kaye, K. Chen, and O. Parcollet, *Phys. Rev. B* **105**, 235115 (2022).
- [166] in *Collected Papers of L.D. Landau*, edited by D. TER HAAR (Pergamon, 1965) pp. 723–730.
- [167] K. A. Hallberg, *Phys. Rev. B* **52**, R9827 (1995).
- [168] D. J. García, K. Hallberg, and M. J. Rozenberg, *Phys. Rev. Lett.* **93**, 246403 (2004).
- [169] S. Nishimoto and E. Jeckelmann, *Journal of Physics: Condensed Matter* **16**, 613 (2004).
- [170] S. Nishimoto, T. Pruschke, and R. M. Noack, *Journal of Physics: Condensed Matter* **18**, 981 (2006).
- [171] S. R. White and A. E. Feiguin, *Phys. Rev. Lett.* **93**, 076401 (2004).
- [172] X. Cao, Y. Lu, P. Hansmann, and M. W. Haverkort, *Physical Review B* **104**, 10.1103/physrevb.104.115119 (2021).
- [173] D. Bauernfeind, X. Cao, E. M. Stoudenmire, and O. Parcollet, *Phys. Rev. B* **105**, 195107 (2022).
- [174] N. Shenvi, J. R. Schmidt, S. T. Edwards, and J. C. Tully, *Phys. Rev. A* **78**, 022502 (2008).
- [175] I. de Vega, U. Schollwöck, and F. A. e. Wolf, *Physical Review B* **92**, 155126 (2015).
- [176] C. Gramsch, K. Balzer, M. Eckstein, and M. Kollar, *Phys. Rev. B* **88**, 235106 (2013).
- [177] X. Dong, E. Gull, and H. U. R. Strand, *Physical Review B* **106**, 10.1103/physrevb.106.125153 (2022).
- [178] R. Bulla, A. C. Hewson, and T. Pruschke, *Journal of Physics: Condensed Matter* **10**, 8365 (1998).
- [179] F. B. Kugler, *Physical Review B* **105**, 10.1103/physrevb.105.245132 (2022).

- [180] M. Karski, C. Raas, and G. S. Uhrig, *Physical Review B* **77**, 10.1103/physrevb.77.075116 (2008).
- [181] C. Krumnow, L. Veis, O. Legeza, and J. Eisert, *Phys. Rev. Lett.* **117**, 210402 (2016).
- [182] R. Schade and P. E. Blöchl, *Physical Review B* **97**, 10.1103/physrevb.97.245131 (2018).
- [183] R. Schade, *New methods for the ab-initio simulation of correlated systems* (2019).
- [184] A.-K. Wu, M. T. Fishman, J. H. Pixley, and E. M. Stoudenmire, *Disentangling interacting systems with fermionic gaussian circuits: Application to the single impurity anderson model* (2022), arXiv:2212.09798 .
- [185] G. Barcza, O. Legeza, K. H. Marti, and M. Reiher, *Physical Review A* **83**, 10.1103/physreva.83.012508 (2011).
- [186] J. Rissler, R. M. Noack, and S. R. White, *Chemical Physics* **323**, 519 (2006).
- [187] H. D. Simon, *Linear Algebra and its Applications* **61**, 101 (1984).
- [188] L. Ding, S. Mardazad, S. Das, S. Szalay, U. Schollwöck, Z. Zimborás, and C. Schilling, *Journal of Chemical Theory and Computation* **17**, 79–95 (2020).
- [189] D. Pekker and B. K. Clark, *Physical Review B* **95**, 10.1103/physrevb.95.035116 (2017).
- [190] Y. Takahashi and H. Umezawa, *International Journal of Modern Physics B* **10**, 1755 (1996), <https://doi.org/10.1142/S0217979296000817> .
- [191] I. de Vega and M.-C. Bañuls, *Phys. Rev. A* **92**, 052116 (2015).
- [192] L. Kohn and G. E. Santoro, *Physical Review B* **104**, 10.1103/physrevb.104.014303 (2021).
- [193] Y. Lu, X. Cao, P. Hansmann, and M. W. Haverkort, *Phys. Rev. B* **100**, 115134 (2019).
- [194] P.-O. Löwdin and H. Shull, *Phys. Rev.* **101**, 1730 (1956).
- [195] R. Schade, *New methods for the ab-initio simulation of correlated systems*, Ph.D. thesis, Georg-August-Universität Göttingen (2019).
- [196] E. M. Stoudenmire and S. R. White, *New Journal of Physics* **12**, 055026 (2010).
- [197] C. Karrasch, J. H. Bardarson, and J. E. Moore, *New Journal of Physics* **15**, 083031 (2013).
- [198] T. Barthel, *New Journal of Physics* **15**, 073010 (2013).

- [199] D. Bauernfeind and M. Aichhorn, *SciPost Phys.* **8**, 024 (2020).
- [200] D. Bauernfeind, R. Triebl, M. Zingl, M. Aichhorn, and H. G. Evertz, *Physical Review B* **97**, 10.1103/physrevb.97.115156 (2018).
- [201] A. Hampel, personal communication.
- [202] M. Fishman, S. R. White, and E. M. Stoudenmire, *SciPost Phys. Codebases* , 4 (2022).
- [203] M. Fishman, S. R. White, and E. M. Stoudenmire, *SciPost Phys. Codebases* , 4 (2022).
- [204] A. Baiardi and M. Reiher, *Journal of Chemical Theory and Computation* **15**, 3481–3498 (2019).
- [205] A. Holzner, A. Weichselbaum, I. P. McCulloch, U. Schollwöck, and J. von Delft, *Phys. Rev. B* **83**, 195115 (2011).
- [206] F. A. Wolf, J. A. Justiniano, I. P. McCulloch, and U. Schollwöck, *Phys. Rev. B* **91**, 115144 (2015).
- [207] K. Guther, W. Dobrautz, O. Gunnarsson, and A. Alavi, *Phys. Rev. Lett.* **121**, 056401 (2018).
- [208] X. Cao, Y. Lu, E. M. Stoudenmire, and O. Parcollet, Dynamical correlation functions from complex time evolution (2023), arXiv:2311.10909 .
- [209] D. N. Page, *Phys. Rev. Lett.* **71**, 1291 (1993).
- [210] A. Weichselbaum and J. von Delft, *Phys. Rev. Lett.* **99**, 076402 (2007).
- [211] R. Peters, T. Pruschke, and F. B. Anders, *Phys. Rev. B* **74**, 245114 (2006).
- [212] A. Weichselbaum, *Phys. Rev. B* **86**, 245124 (2012).
- [213] S.-S. B. Lee, J. von Delft, and A. Weichselbaum, *Phys. Rev. Lett.* **119**, 236402 (2017).
- [214] S.-S. B. Lee and A. Weichselbaum, *Phys. Rev. B* **94**, 235127 (2016).
- [215] A. Weichselbaum, *Annals of Physics* **327**, 2972 (2012).
- [216] M. Jarrell and J. E. Gubernatis, *Physics Reports* **269**, 133 (1996).
- [217] K. M. Stadler, Z. P. Yin, J. von Delft, G. Kotliar, and A. Weichselbaum, *Phys. Rev. Lett.* **115**, 136401 (2015).
- [218] F. B. Kugler, S.-S. B. Lee, A. Weichselbaum, G. Kotliar, and J. von Delft, *Phys. Rev. B* **100**, 115159 (2019).

- [219] A. K. Mitchell, M. R. Galpin, S. Wilson-Fletcher, D. E. Logan, and R. Bulla, Phys. Rev. B **89**, 121105 (2014).
- [220] K. M. Stadler, A. K. Mitchell, J. von Delft, and A. Weichselbaum, Phys. Rev. B **93**, 235101 (2016).
- [221] F. B. Kugler and G. Kotliar, Phys. Rev. Lett. **129**, 096403 (2022).
- [222] J. Kondo, Progress of Theoretical Physics **32**, 37 (1964), <https://academic.oup.com/ptp/article-pdf/32/1/37/5193092/32-1-37.pdf> .
- [223] K. Andres, J. E. Graebner, and H. R. Ott, Phys. Rev. Lett. **35**, 1779 (1975).
- [224] F. Steglich, J. Aarts, C. D. Bredl, W. Lieke, D. Meschede, W. Franz, and H. Schäfer, Phys. Rev. Lett. **43**, 1892 (1979).
- [225] N. D. Mathur, F. M. Grosche, S. R. Julian, I. R. Walker, D. M. Freye, R. K. W. Haselwimmer, and G. G. Lonzarich, Nature **394**, 39 (1998).
- [226] H. v. Löhneysen, T. Pietrus, G. Portisch, H. G. Schlager, A. Schröder, M. Sieck, and T. Trappmann, Phys. Rev. Lett. **72**, 3262 (1994).
- [227] R. Ballou, E. Lelièvre-Berna, and B. Fåk, Phys. Rev. Lett. **76**, 2125 (1996).
- [228] V. I. Anisimov, M. A. Korotin, M. Zöfl, T. Pruschke, K. Le Hur, and T. M. Rice, Phys. Rev. Lett. **83**, 364 (1999).
- [229] C. Urano, M. Nohara, S. Kondo, F. Sakai, H. Takagi, T. Shiraki, and T. Okubo, Phys. Rev. Lett. **85**, 1052 (2000).
- [230] S. Kondo, C. Urano, Y. Kurihara, M. Nohara, and H. Takagi, *Magnetic Excitations in Strongly Correlated Electrons: From the Geometrically Frustrated Antiferromagnets ZnV_2O_4 and $ZnCr_2O_4$ to the Heavy-Mass Fermi Liquid LiV_2O_4* , Vol. 69 (2000) pp. 139–143.
- [231] Y. Shimizu, H. Takeda, M. Tanaka, M. Itoh, S. Niitaka, and H. Takagi, Nature Communications **3**, 981 (2012).
- [232] K. Tomiyasu, K. Iwasa, H. Ueda, S. Niitaka, H. Takagi, S. Ohira-Kawamura, T. Kikuchi, Y. Inamura, K. Nakajima, and K. Yamada, Phys. Rev. Lett. **113**, 236402 (2014).
- [233] B.-C. Gong, H.-C. Yang, K. Jin, K. Liu, and Z.-Y. Lu, Chinese Physics B **29**, 077508 (2020).
- [234] G. Cao, S. McCall, M. Shepard, J. E. Crow, and R. P. Guertin, Phys. Rev. B **56**, 321 (1997).

- [235] Y. Maeno, H. Hashimoto, K. Yoshida, S. Nishizaki, T. Fujita, J. G. Bednorz, and F. Lichtenberg, *Nature* **372**, 532 (1994).
- [236] Y. Maeno, T. M. Rice, and M. Sigrist, *Physics Today* **54**, 42 (2001), https://pubs.aip.org/physicstoday/article-pdf/54/1/42/11109470/42_1_online.pdf .
- [237] H. A. Krug von Nidda, R. Bulla, N. Büttgen, M. Heinrich, and A. Loidl, *The European Physical Journal B - Condensed Matter and Complex Systems* **34**, 399 (2003).
- [238] R. Arita, K. Held, A. V. Lukoyanov, and V. I. Anisimov, *Phys. Rev. Lett.* **98**, 166402 (2007).
- [239] M. Feldbacher, K. Held, and F. F. Assaad, *Phys. Rev. Lett.* **93**, 136405 (2004).
- [240] R. Arita and K. Held, *Phys. Rev. B* **72**, 201102 (2005).
- [241] M. Crispino, P. V. Arribi, A. Shukla, F. Hardy, A.-A. Haghighirad, T. Wolf, R. Heid, C. Meingast, T. Gorni, A. Avella, and L. de' Medici, Paradigm for finding d-electron heavy fermions: the case of cr-doped csfe_2as_2 (2023), arXiv:2312.06511 [cond-mat.str-el] .
- [242] K. Nakamura, Y. Yoshimoto, Y. Nomura, T. Tadano, M. Kawamura, T. Kosugi, K. Yoshimi, T. Misawa, and Y. Motoyama, *Computer Physics Communications* **261**, 107781 (2021).
- [243] O. Parcollet, M. Ferrero, T. Ayrál, H. Hafermann, I. Krivenko, L. Messio, and P. Seth, *Computer Physics Communications* **196**, 398 (2015).
- [244] M. Aichhorn, L. Pourovskii, P. Seth, V. Vildosola, M. Zingl, O. E. Peil, X. Deng, J. Mravlje, G. J. Krabberger, C. Martins, M. Ferrero, and O. Parcollet, *Computer Physics Communications* **204**, 200 (2016).
- [245] P. Seth, I. Krivenko, M. Ferrero, and O. Parcollet, *Computer Physics Communications* **200**, 274 (2016).
- [246] M. E. Merkel, A. Carta, S. Beck, and A. Hampel, *The Journal of Open Source Software* **7**, 4623 (2022).
- [247] F. Verstraete, J. J. García-Ripoll, and J. I. Cirac, *Phys. Rev. Lett.* **93**, 207204 (2004).
- [248] J. Haegeman, J. I. Cirac, T. J. Osborne, I. Pižorn, H. Verschelde, and F. Verstraete, *Phys. Rev. Lett.* **107**, 070601 (2011).
- [249] A. V. Mahajan, R. Sala, E. Lee, F. Borsa, S. Kondo, and D. C. Johnston, *Phys. Rev. B* **57**, 8890 (1998).

-
- [250] X. Cao, E. M. Stoudenmire, and O. Parcollet, Finite temperature minimal entangled typical thermal states impurity solver (2023), arXiv:2312.13668 [cond-mat.str-el] .
- [251] O. Parcollet, M. Ferrero, T. Ayrat, H. Hafermann, I. Krivenko, L. Messio, and P. Seth, *Computer Physics Communications* **196**, 398 (2015).
- [252] I. A. Nekrasov, Z. V. Pchelkina, G. Keller, T. Pruschke, K. Held, A. Krimmel, D. Vollhardt, and V. I. Anisimov, *Phys. Rev. B* **67**, 085111 (2003).
- [253] J. Matsuno, A. Fujimori, and L. F. Mattheiss, *Phys. Rev. B* **60**, 1607 (1999).
- [254] P. Werner, E. Gull, M. Troyer, and A. J. Millis, *Phys. Rev. Lett.* **101**, 166405 (2008).
- [255] A. Kowalski, A. Hausoel, M. Wallerberger, P. Gunacker, and G. Sangiovanni, *Phys. Rev. B* **99**, 155112 (2019).
- [256] C. Hubig, I. P. McCulloch, and U. Schollwöck, *Phys. Rev. B* **95**, 035129 (2017).
- [257] S. Paeckel, T. Köhler, and S. R. Manmana, *SciPost Phys.* **3**, 035 (2017).

Acknowledgements

This thesis was made possible by the support of many. In particular, I would like to thank

- **Ulrich Schollwöck** for providing valuable guidance throughout my academic endeavour. Thank you for always finding time, and supporting me all the way. I truly enjoyed our conversations.
- **Sebastian Paeckel** for bringing a spark of chaos into the German officialdom. It is a great pleasure to work with you. Mostly. But please, do not come around to watch a football match anymore.
- **Max Bramberger** for being a fantastic colleague and friend. Creating OTIS with you was not only the cornerstone of this thesis; it will forever remain a cherished memory.
- **Sam Mardazad** for being a fantastic colleague and friend. You helped me immensely in mastering the depths of SYTEN and finding appreciation for a good book and a nice glass of whiskey.
- **Mattia Moroder** for being a fantastic colleague and friend. You have been my first master student, and it brings me immense joy to watch you complete your PhD alongside me.
- I want to thank the Simons Foundation for their hospitality. In particular, I want to thank **Olivier Parcollet**, **Alexander Hampel**, **Xiaodong Cao**, **Fabian Kugler**, **Antoine Georges** and **Andy Millis**. The stay at CCQ rekindled my love for science and all of you made it a wonderful experience.
- **Claudius Hubig** for introducing me to the topic of tensor networks and for the incredible work that went into SYTEN. Without it, many aspects of this thesis would have been significantly harder, or frankly impossible.
- **Nils-Oliver Linden** for introducing me to DMFT and thereby inspiring this work. You have been a fantastic supervisor.
- I want to thank all my former bachelor and master students for all their hard work and effort. It was a pleasure to work with every single one of you.

- **Cordula Weber** for abating the burden of bureaucracy. I am immensely grateful for the wonderful work you did, even though I am probably not even aware of half of it.
- I want to thank my friends for bearing with me throughout this period. It was too often that I thought I did not have the time to join you. How rather foolish of me. I am very much looking forward to changing this.
- My parents for always supporting me during this journey. I am immensely grateful for all your efforts, especially those focused on keeping the mood up high.
- **Iris** for, well, everything. We started our PhD's together and I could not imagine myself finishing it without you. You always had my back. You were right beside me during the best of times, during the worst of times. Thank you.