Massively Parallel Simulations of Correlated Materials

X.-J. Zhang, N. Rihaisamani, E. Adibi, E. Koch, E. Pavarini

published in

NIC Symposium 2022
M. Müller, Ch. Peter, A. Trautmann (Editors)

Forschungszentrum Jülich GmbH,
John von Neumann Institute for Computing (NIC),
Schriften des Forschungszentrums Jülich, NIC Series, Vol. 51,
ISBN 978-3-95806-646-5, pp. 289.
http://hdl.handle.net/2128/31840

Massively Parallel Simulations of Correlated Materials

Xue-Jing Zhang, Neda Rihaisamani, Elaheh Adibi, Erik Koch, and Eva Pavarini

Institute for Advanced Simulation, Forschungszentrum Jülich, 52425 Jülich, Germany
E-mail: {x.zhang, n.rihaisamani, e.adibi, e.koch, e.pavarini}@fz-juelich.de

The realistic description of strongly-correlated systems is one of the great challenges in condensed-matter physics. In the last decade, impressive progress has been achieved by combining new algorithms with the power of massively-parallel supercomputers. Recent developments made it possible to systematically explore not only spectral functions but also static and dynamical responses. Furthermore, new categories of systems, such as those with important spin-orbit interaction effects, including topologically non-trivial materials, became accessible. Here we briefly describe the numerical approach behind this success.

1 Introduction

Electronic-structure theory is the basis of modern technologies such as electronics and computing. Electronic properties of materials are determined by quantum mechanics. Thus, in principle, by solving the Schrödinger equation, we should be able to predict the properties of real materials, or even design new ones with superior qualities. In practice, unfortunately, solving this equation is not easy at all. The essential complication comes from the inherent quantum many-body nature of the problem. As a result, a brute-force solution is impossible, except in the simplest cases. As an illustration let us consider a single atom of iron. Having 26 electrons, its wave function is a function of 26 times 3 coordinates. Neglecting spin, already an extremely crude representation of this function at merely 10 values of each variable would thus require storage of 10^{78} numbers. Even after reducing this number by exploiting symmetries, there is simply not enough matter available in our galaxy for building the required memory.

Given this example, electronic-structure theory seems a hopeless enterprise. Nevertheless, it is a thriving discipline. This is largely due to density-functional theory (DFT). In practice, thanks to the Kohn-Sham (KS) picture, this approach drastically simplifies the many-body problem by assuming that the electrons retain their individuality and experience the other electrons via a static mean field. In this picture electrons occupy states that extend over the whole crystal, forming the band structure of the material.

For many important classes of materials a density-functional description fails even qualitatively, however. Striking effects like the breakdown of the Fermi-liquid picture at the Mott metal-insulator transition, heavy Fermion behaviour, exotic one-dimensional Luttinger phases, or high-temperature superconductivity cannot be addressed by such a simple approach. All these materials are called *strongly-correlated* systems. This name originates from the fact that, because of the strong repulsion between the electrons, the latter lose their individuality, and the single-particle picture breaks down. Furthermore, because of the strength of the interaction, non-perturbative many-body techniques have to be used, so that very powerful computers are essential for reliable calculations. And still, calculations are restricted to quite small model systems. This means that the full Hamiltonian of a crystal has to be approximated by a small lattice Hamiltonian, which describes only (few of)

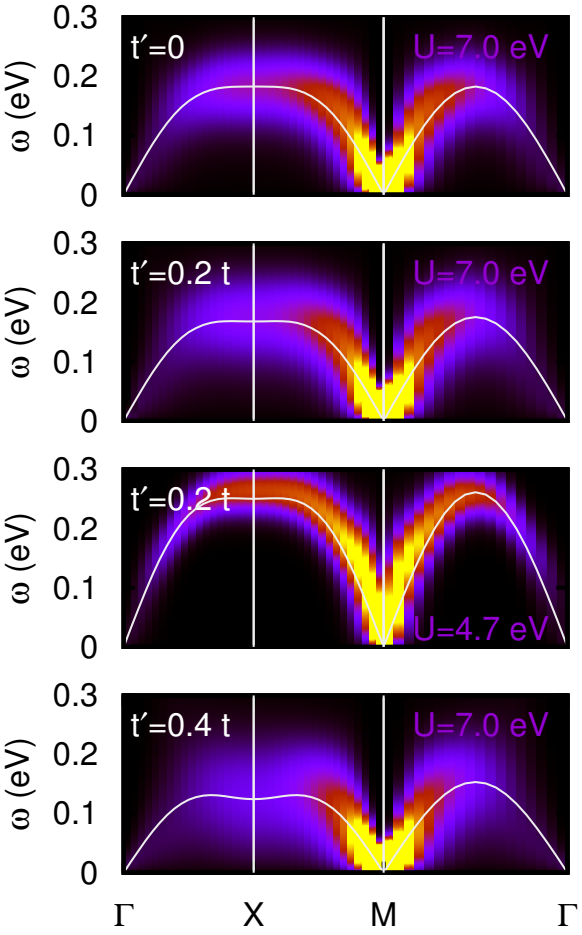

Figure 1. Dynamical response function yielding bosonic Goldstone modes for a model simulating high-temperature superconducting cuprates (reproduced from Ref. 14).

the strongly-correlated electrons. All other electrons have to be included in the calculation in an average way[1].

The modern approach to solving the many-body problem is dynamical mean-field theory (DMFT)[2–7]. It reduces the lattice Hamiltonian to a correlated impurity embedded in a self-consistent dynamical medium, which mimics the other lattice sites. This approximation simplifies the problem significantly. Still, a complicated quantum-impurity problem remains, which has to be solved. One of the most flexible methods to this aim is quantum Monte Carlo (QMC)[8–12]. Thanks to DMFT, it was possible, for the first time, to understand the physics of the Mott transition. In Mott insulators the electronic band structure loses its meaning. Instead, physics becomes more local and it is more appropriate to think about the electrons as occupying atomic-like orbitals.

In this strongly-correlated regime we find a number of fascinating ordering phenom-

ena. Most well known is anti-ferromagnetism, where spins on neighbouring lattice sites point in opposite directions. When there are many correlated orbitals, a similar ordered phase can exist: occupied orbitals on neighbouring sites point in different directions. This directionality can give rise to highly anisotropic transport properties. Coupling of spin- and orbital degrees of freedom can make transport properties strongly dependent on magnetic field. Such a mechanism is believed to be the basis of the colossal magneto-resistance effect (CMR). Like the giant magneto-resistance (GMR), this effect, once fully understood, holds the promise of, e.g., another vast increase in hard-disk capacity.

The realistic description of most interesting systems requires that both spin and orbital degrees of freedom are accounted for. This enlarges the size of the quantum impurity problem; the computational time scales with a power (typically larger than one) of the number of the degrees of freedom; which power exactly depending on the specific algorithm and problem. A further doubling of size comes from the spin-orbit interaction, which is key for the description of many families of compounds, in particular topologically non-trivial systems. Finally, complicated spatial patterns cannot be described by a single-site approach such as DMFT, which assumes that all lattice sites are equivalent. In order to add the required spatial degrees of freedom the single impurity of DMFT has to be replaced by a cluster of sites. This approach is accordingly called cluster DMFT (CDMFT). Unfortunately, treating a large cluster instead of a single site increases the already high computational cost of a calculation even further.

With an efficient parallelisation of the QMC solver and of the DMFT self-consistency loop, we can however exploit the spectacular increase in performance offered by massively parallel machines of the class of JUWELS. Thanks to advances in algorithms and supercomputers, it is now possible to reach experimental temperatures, where, before massively-parallel supercomputers, calculations of low temperature physics were limited to temperatures of about 1000 K. In addition, for lack of computer time, uncontrolled approximations had to be introduced in the model Hamiltonians. Now it is instead possible to check these approximations by explicit calculations. In short, calculations are becoming significantly more reliable and consequently gain predictive power. In most recent years, it became possible to calculate systematically statical and dynamical response functions, obtain bosonic excitation spectra (an example is shown in Fig. 1), and study spin-orbit correlated materials[13–20], included topologically non-trivial systems.

2 Method

2.1 DMFT

The Hubbard Hamiltonian is the simplest model for the description of the Mott metal-insulator transition. In the tight-binding approximation it becomes

$$\hat{H} = \varepsilon_d \sum_{\sigma i} \hat{n}_{i\sigma} - t \sum_{\sigma \langle ii' \rangle} c^\dagger_{i\sigma} c_{i'\sigma} + U \sum_i \hat{n}_{i\uparrow} \hat{n}_{i\downarrow}, \tag{1}$$

where $c^\dagger_{i\sigma}$ ($c_{i\sigma}$) creates (destroys) an electron at site i and with spin σ, and $\langle ii' \rangle$ is a sum over first neighbours. Here the competing energy scales are the hopping integral, t, and the on-site Coulomb repulsion, U. The on-site energy, ε_d, instead, only modifies the chemical

potential and thus does not affect physical properties. For $U{=}0$, at half-filling, the Hamiltonian in Eq. 1 describes a metallic band. For $t{=}0$ it describes an insulating collection of disconnected atoms. Somewhere in between, at a critical value of t/U, a metal to insulator transition must occur. Here we will discuss the DMFT solution of Eq. 1 and the picture of the metal-insulator transition emerging from it. The first step consists in mapping Eq. 1 into an effective quantum-impurity model, such as the Anderson Hamiltonian

$$\hat{H}^A = \underbrace{\sum_{\mathbf{k}\sigma} \varepsilon^s_{\mathbf{k}} \hat{n}_{\mathbf{k}\sigma}}_{\hat{H}_{\rm bath}} + \underbrace{\sum_{\mathbf{k}\sigma} \left(V^s_{\mathbf{k}} c^\dagger_{\mathbf{k}\sigma} c_{d\sigma} + \text{h.c.} \right)}_{\hat{H}_{\rm hyb}} + \underbrace{\varepsilon_d \sum_\sigma \hat{n}_{d\sigma} + U \hat{n}_{d\uparrow} \hat{n}_{d\downarrow}}_{\hat{H}_{\rm imp}} . \tag{2}$$

In this model the on-site Coulomb repulsion U appears only in the impurity Hamiltonian, $\hat{H}_{\rm imp}$, while the terms $\hat{H}_{\rm bath}$ and $\hat{H}_{\rm hyb}$, describe, respectively, the bath and the bath-impurity hybridisation. In the next step, the quantum-impurity model is solved. This requires non-perturbative numerical methods, such as exact diagonalisation, the numerical renormalisation group, density-matrix renormalisation group or QMC. Here we describe the DMFT self-consistency loop for a QMC quantum-impurity solver. Solving the quantum-impurity model yields the impurity Green function $G^\sigma_{d,d}(i\nu_n)$. From the impurity Dyson equation we can calculate the impurity self-energy

$$\Sigma^\sigma_A(i\nu_n) = \left(G^{0\sigma}_{d,d}(i\nu_n)\right)^{-1} - \left(G^\sigma_{d,d}(i\nu_n)\right)^{-1}. \tag{3}$$

Next, we adopt the local self-energy approximation, i.e., we assume that the self-energy of the Hubbard model equals the impurity self-energy. The local Dyson equation is used once more, this time to calculate the bath Green function $\mathcal{G}^\sigma(i\nu_n)$, which in turn defines a new quantum-impurity model. This procedure is repeated until self-consistency is reached, i.e., the number of electrons is correct and the self-energy does not change anymore (within a given numerical accuracy). It is important to underline that self-consistency is key to the success of DMFT in describing the metal-to-insulator transition. The self-consistency loop is shown schematically in Fig. 2 for the general case.

Although the extension of DMFT to Hubbard models with many orbitals might appear straightforward, in practice it is not. The bottleneck is the solution of the generalised multi-orbital quantum-impurity problem via QMC. Despite being flexible, QMC-based approaches have limitations. These can be classified in two types. First, with increasing the number of degrees of freedom, calculations become very quickly computationally too expensive – how quickly depends on the specific QMC algorithm used and the actual implementation. Thus, going beyond a rather small number of orbitals and reaching the zero-temperature limit is unfeasible in practice. The second type of limitation is more severe. Increasing the number of degrees of freedom leads, eventually, to the infamous sign problem; when this happens, QMC calculations cannot be performed at all. In order to deal with limitations of the first type, it is crucial to restrict QMC calculations to the essential degrees of freedom; furthermore, we should exploit symmetries, develop fast algorithms and use the power of massively parallel supercomputers to reduce the actual computational time. For the second type of problems not a lot can be done; nevertheless, it has been shown that a severe sign problem might appear earlier with some basis choices than with others[12]. Although eventually we cannot escape it, this suggests that the model set-up can be used as a tool to expand the moderate sign-problem zone.

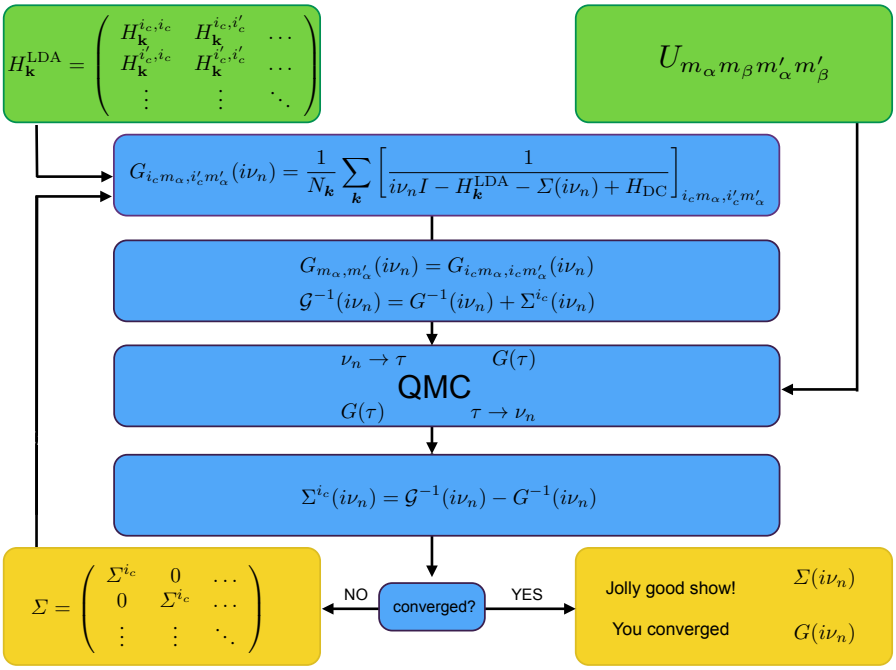

Figure 2. DMFT self-consistency loop. The one-electron Hamiltonian is built in the basis of Bloch states obtained from localised Wannier functions, e.g., in the local-density approximation (LDA); this gives $H^{\mathrm{LDA}}_{\mathbf{k}}$. The set $\{i_c\}$ labels the equivalent correlated sites inside the unit cell. The local Green-function matrix is at first calculated using an initial guess for the self-energy matrix. The bath Green-function matrix is then obtained via the Dyson equation and used to construct an effective quantum-impurity model. The latter is solved via a quantum-impurity solver, here quantum Montecarlo (QMC). This yields the impurity Green-function matrix. Through the Dyson equation the self-energy is then obtained. The procedure is repeated until self-consistency is reached (reproduced from Ref. 7).

2.2 Quantum-Impurity Solvers: Continuous-Time Quantum Monte Carlo

Here we explain the principles of the continuous-time QMC method[8–12], perhaps the most successful among QMC-based quantum-impurity solvers. In particular, we illustrate the hybridisation expansion version of the approach. To make things simpler, we consider the case of a Hubbard model in Eq. 1 with only two sites, $i = 1, 2$, and, correspondingly, a quantum-impurity model of form Eq. 2 also made of two sites, one representing the impurity, labelled with d, and one representing the bath, labelled with s. The first step consists in splitting the quantum-impurity Hamiltonian into bath ($\hat{H}_{\mathrm{bath}}$), hybridisation ($\hat{H}_{\mathrm{hyb}}$) and local ($\hat{H}_{\mathrm{loc}}$) terms

$$\hat{H}^{\mathrm{A}} = \underbrace{\varepsilon_s \sum_\sigma \hat{n}_{s\sigma}}_{\hat{H}_{\mathrm{bath}}} \underbrace{-t \sum_\sigma \left(c^\dagger_{d\sigma} c_{s\sigma} + c^\dagger_{s\sigma} c_{d\sigma} \right)}_{\hat{H}_{\mathrm{hyb}}} + \underbrace{\varepsilon_d \sum_\sigma \hat{n}_{d\sigma} + U \hat{n}_{d\uparrow} \hat{n}_{d\downarrow}}_{\hat{H}_{\mathrm{loc}}} . \qquad (4)$$

Next, we write the partition function Z as a perturbation series in the hybridisation. To this end, we define $\hat{H}_0 = \hat{H}_{\rm bath} + \hat{H}_{\rm loc}$ and rewrite the partition function as

$$Z = \mathrm{Tr}\left(e^{-\beta(\hat{H}_0-\mu\hat{N})}\hat{V}(\beta)\right), \tag{5}$$

where the operator $\hat{V}(\beta)$ is given by

$$\hat{V}(\beta) = e^{\beta(\hat{H}_0-\mu\hat{N})}\, e^{-\beta(\hat{H}_0+\hat{H}_{\rm hyb}-\mu\hat{N})} = \sum_m \underbrace{\int_0^\beta d\tau_1 \cdots \int_{\tau_{m-1}}^\beta d\tau_m}_{\int d\tau^m} \underbrace{(-1)^m \prod_{l=m}^1 \hat{H}_{\rm hyb}(\tau_l)}_{\hat{O}^m(\tau)}, \tag{6}$$

and

$$\hat{H}_{\rm hyb}(\tau_l) = e^{\tau_l(\hat{H}_0-\mu\hat{N})}\, \hat{H}_{\rm hyb}\, e^{-\tau_l(\hat{H}_0-\mu\hat{N})}. \tag{7}$$

In this expansion, the only terms that contribute to the trace are even order ones ($m = 2k$) and they are products of impurity (d) and bath (s) creator-annihilator pairs. Thus

$$\int d\tau^{2k} \longrightarrow \int d\tau^k \int d\bar{\tau}^k \quad \text{and} \qquad \hat{O}^{2k}(\tau) \longrightarrow \sum_{\sigma,\bar{\sigma}} \hat{O}^{2k}_{\sigma,\bar{\sigma}}(\tau,\bar{\tau}), \tag{8}$$

where

$$\hat{O}^{2k}_{\sigma,\bar{\sigma}}(\tau,\bar{\tau}) = (t)^{2k} \prod_{i=1}^k \left(c^\dagger_{d\bar{\sigma}_i}(\bar{\tau}_i) c_{s\bar{\sigma}_i}(\bar{\tau}_i) c^\dagger_{s\sigma_i}(\tau_i) c_{d\sigma_i}(\tau_i)\right). \tag{9}$$

The vector $\sigma = (\sigma_1, \sigma_2, ..., \sigma_k)$ gives the spins $\{\sigma_i\}$ associated with the k impurity annihilators at imaginary times $\{\tau_i\}$, while $\bar{\sigma} = (\bar{\sigma}_1, \bar{\sigma}_2, ..., \bar{\sigma}_k)$ gives the spins $\{\bar{\sigma}_i\}$ associated with the k impurity creators at imaginary times $\{\bar{\tau}_i\}$. It follows that the local and bath traces can be decoupled and the partition function can be rewritten as

$$\frac{Z}{Z_{\rm bath}} = \sum_k \int d\tau^k \int d\bar{\tau}^k \sum_{\sigma,\bar{\sigma}} d^k_{\bar{\sigma},\sigma}(\tau,\bar{\tau})\, \mathbf{t}^{\mathbf{k}}_{\sigma,\bar{\sigma}}(\tau,\bar{\tau}), \tag{10}$$

$$d^k_{\bar{\sigma},\sigma}(\tau,\bar{\tau}) = \frac{t^{2k}}{Z_{\rm bath}}\, \mathrm{Tr}_{\rm bath}\left(e^{-\beta(\hat{H}_{\rm bath}-\mu\hat{N}_s)} \mathcal{T} \prod_{i=k}^1 c^\dagger_{s\sigma_i}(\tau_i) c_{s\bar{\sigma}_i}(\bar{\tau}_i)\right), \tag{11}$$

$$t^k_{\sigma,\bar{\sigma}}(\tau,\bar{\tau}) = \mathrm{Tr}_{\rm loc}\left(e^{-\beta(\hat{H}_{\rm loc}-\mu\hat{N}_d)} \mathcal{T} \prod_{i=k}^1 c_{d\sigma_i}(\tau_i) c^\dagger_{d\bar{\sigma}_i}(\bar{\tau}_i)\right), \tag{12}$$

where $Z_{\rm bath} = 1 + 2e^{-\beta(\varepsilon_s-\mu)} + e^{-2\beta(\varepsilon_s-\mu)}$. The trace involving only bath operators is simple to calculate, since $\hat{H}_{\rm bath}$ describes an independent-electron problem for which Wick's theorem holds. It is given by the determinant

$$d^k_{\bar{\sigma},\sigma}(\tau,\bar{\tau}) = \det\left(\mathcal{F}^k_{\bar{\sigma},\sigma}(\tau,\bar{\tau})\right), \tag{13}$$

of the $k\times k$ non-interacting hybridisation-function matrix, with elements

$$\left(\mathcal{F}^k_{\bar{\sigma},\sigma}(\tau,\bar{\tau})\right)_{i',i} = \mathcal{F}^0_{\bar{\sigma}_{i'},\sigma_i}(\bar{\tau}_{i'}-\tau_i), \tag{14}$$

where

$$\mathcal{F}^0_{\bar{\sigma},\sigma}(\tau) = \delta_{\bar{\sigma},\sigma} \frac{t^2}{1+e^{-\beta(\varepsilon_s-\mu)}} \times \begin{cases} -e^{-\tau(\varepsilon_s-\mu)} & \tau > 0, \\ +e^{-(\beta+\tau)(\varepsilon_s-\mu)} & \tau < 0. \end{cases} \tag{15}$$

The calculation of the local trace is in general more complicated. In the case discussed here, the Hamiltonian does not flip spins. Thus only terms with an equal number of creation and annihilation operators *per spin* contribute to the local trace, and we can express the partition function in expansion orders per spin, k_σ. This yields[12]

$$\frac{Z}{Z_{\rm bath}} = \left(\prod_\sigma \sum_{k_\sigma=0}^{\infty} \int d\tau_\sigma^{k_\sigma} \int d\bar{\tau}_\sigma^{k_\sigma} \right) d^k_{\bar{\sigma},\sigma}(\tau,\bar{\tau}) t^k_{\sigma,\bar{\sigma}}(\tau,\bar{\tau}) \tag{16}$$

where the vectors $\sigma = (\sigma_\uparrow, \sigma_\downarrow)$ and $\bar{\sigma} = (\bar{\sigma}_\uparrow, \bar{\sigma}_\downarrow)$ have $(k_\uparrow, k_\downarrow)$ components, and for each k_σ component $\sigma_i = \bar{\sigma}_i = \sigma$. Thus

$$t^k_{\sigma,\bar{\sigma}}(\tau,\bar{\tau}) = {\rm Tr}_{\rm loc} \left(e^{-\beta(\hat{H}_{\rm loc}-\mu\hat{N}_d)} \mathcal{T} \prod\nolimits_\sigma \prod\nolimits_{i=k_\sigma}^{1} c_{d\sigma}(\tau_{\sigma_i}) c^\dagger_{d\sigma}(\bar{\tau}_{\bar{\sigma}_i}) \right). \tag{17}$$

The latter can be calculated analytically. To do this, first we parametrise all configurations for a given spin via a timeline $[0,\beta)$ plus a number of creator/annihilator pairs which define segments on the timeline. At zeroth order two possible configurations exist per spin, an empty timeline, which corresponds to the vacuum state $|0\rangle$, and a full timeline, which corresponds to the state $c^\dagger_{d\sigma}|0\rangle$. A given configuration yields, at order $k = k_\uparrow + k_\downarrow$

$$t^k_{\sigma,\bar{\sigma}}(\tau,\bar{\tau}) = \left(\prod_\sigma s_\sigma^{k_\sigma} \right) e^{-\sum_{\sigma\sigma'}((\varepsilon_d-\mu)\delta_{\sigma\sigma'} + \frac{U}{2}(1-\delta_{\sigma,\sigma'}))l_{\sigma,\sigma'}}, \tag{18}$$

where $l_{\sigma,\sigma'}$ is the length of the overlap of the τ segments for spins σ and σ', respectively, while $s_\sigma = {\rm sign}(\tau_{\sigma_1} - \bar{\tau}_{\sigma_1})$ is the fermionic sign. Possible configurations at order $k = 0, 1, 2$ are shown in Fig. 3. Taking all k values into account, the partition function can be expressed as the sum over all configurations $\{c\}$, i.e., in short

$$Z = \sum_c w_c = \sum_c |w_c| \ {\rm sign}\ w_c. \tag{19}$$

In a compact form, we can write $w_c = d\tau_c \ d_c \ t_c$ where $d\tau_c = \prod_\sigma \prod_i^{k_\sigma} d\tau_{\sigma_i} d\bar{\tau}_{\bar{\sigma}_i}$, and d_c and t_c are the bath and local traces for the configuration c. This expression of the partition function shows that we can interpret $|w_c|$ as the sampling weight of configuration c. A generic observable $\hat{O}$ can then be obtained as the Monte Carlo average on a finite number of configurations N_c

$$\langle \hat{O} \rangle = \frac{\sum_c \langle \hat{O} \rangle_c |w_c| \ {\rm sign}\ w_c}{\sum_c |w_c| \ {\rm sign}\ w_c} \approx \frac{\frac{1}{N_c} \sum_c^{N_c} \langle \hat{O} \rangle_c \ {\rm sign}\ w_c}{\frac{1}{N_c} \sum_c {\rm sign}\ w_c}. \tag{20}$$

The term $\frac{1}{N_c} \sum_c {\rm sign}\ w_c$ in the denominator is the average fermionic sign. When this is small, much longer runs are required to obtain data of the same quality; eventually the computational time can become so long that the calculation is unfeasible – in these cases we have a sign problem. In practice, the QMC simulation starts from a random

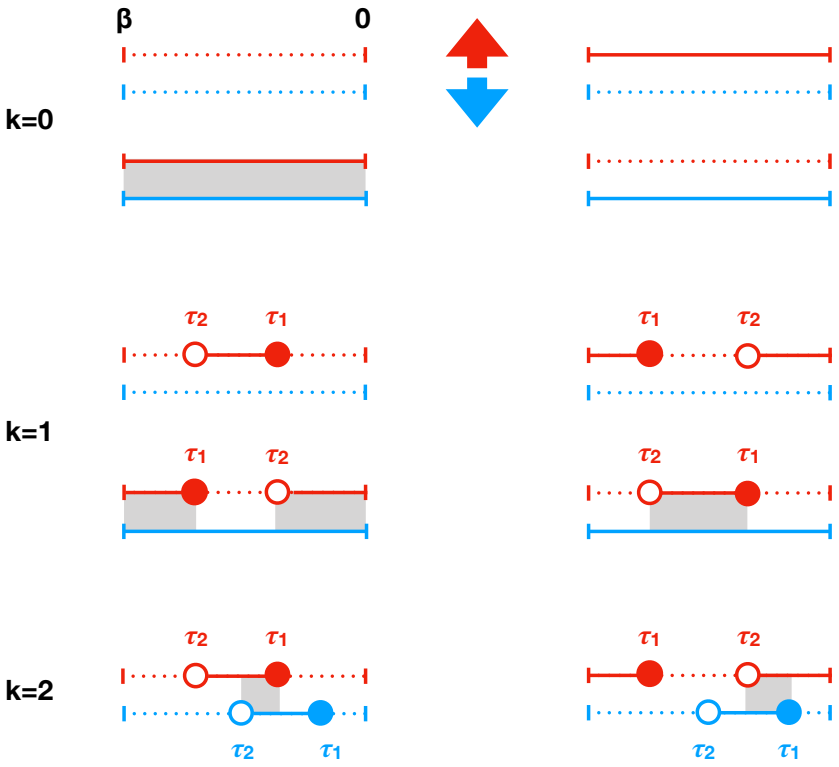

Figure 3. Representative configurations contributing to the local trace at zeroth, first and second order (reproduced from Ref. 7). The timelines for spin up are red and those for spin down are blue. The filled circles correspond to the insertion of a creator (time τ_1), and the empty circles to the insertion of an annihilator (time τ_2). Dotted lines represent the vacuum state for a given spin, full lines the occupied state. The grey boxes indicate the regions in which $l_{\uparrow,\downarrow} \neq 0$.

configuration c. Next we propose an update $c \to c'$. Within the Metropolis algorithm, the acceptance ratio is

$$R_{c\to c'} = \min\left(1, \frac{p_{c'\to c}}{p_{c\to c'}} \frac{|w_{c'}|}{|w_c|}\right), \tag{21}$$

where $p_{c\to c'}$ is the proposal probability for the update $c \to c'$. In the approach described here, known as *segment solver*, the basic updates are addition and removal of segments, antisegments (segments winding over the borders of the timeline, see Fig. 3), or complete lines. As example, let us consider the insertion of a segment for spin σ. A segment is made by a creator and an annihilator. The creator is inserted at time $\tau_{\rm in}$; the move is rejected if $\tau_{\rm in}$ is in a region where a segment exists. If created, the segment can have at most length $l_{\rm max}$, given by the distance between $\tau_{\rm in}$ and the time at which the next creator is, hence

$$p_{c\to c'} = \frac{d\bar{\tau}}{\beta} \frac{d\tau}{l_{\rm max}} . \tag{22}$$

The proposal probability of the reverse move (removing a segment) is instead given by the inverse of the number of existing segments

$$p_{c' \to c} = \frac{1}{k_\sigma + 1} \,. \tag{23}$$

The acceptance ratio for the insertion of a segment becomes then

$$R_{c \to c'} = \min\left(1, \frac{\beta l_{max}}{k_\sigma + 1} \left| \frac{d_{c'}}{d_c} \frac{t_{c'}}{t_c} \right| \right). \tag{24}$$

In multi-orbital and multi-site systems the updates are more complex due to the coupling between flavours, but the general principle holds. Dynamical responses require sampling three independent imaginary times, sizably increasing the computational cost.

3 Concluding Remarks

We have described the key steps behind numerical simulation of correlated electrons in materials. In recent years the ability of calculating systematically static and dynamical response functions (see Fig. 1) and to study spin-orbit materials (see Fig. 4), included topologically non-trivial matter, has opened new horizons in understanding correlated systems. This is not only key for basic physics but also for practical applications. Thanks to their high sensitivity to external parameters such as pressure, temperature or magnetic/electric fields, correlated systems hold the promise to play a key role in future electronics and information technology. Theoretically the challenge is to solve systematically ever larger models, in order to eventually reach predictive power. Experience so far tells us that ultimately this can only be achieved with the help of large-scale simulations on novel generations of massively parallel supercomputers.

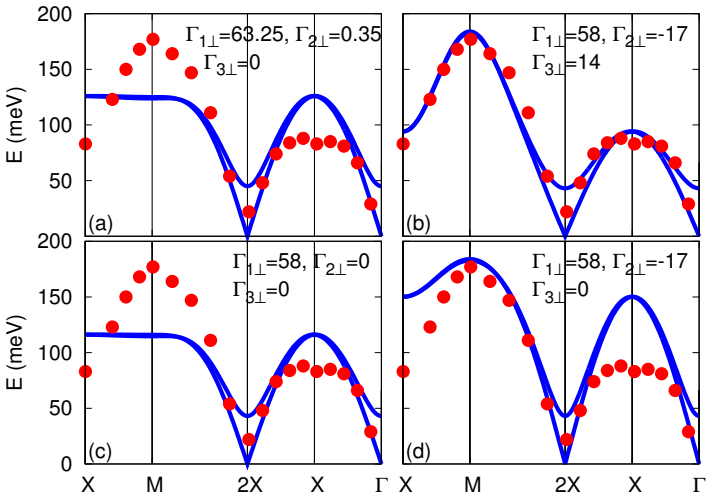

Figure 4. Spin waves in the iridates (reproduced from Ref. 13).

Acknowledgements

Initial code development and benchmark computations were performed with a grant of computer time provided by the VSR/JARA of the Forschungszentrum Jülich. The authors gratefully acknowledge the Gauss Centre for Supercomputing e.V. (`www.gauss-centre.eu`) for funding this project by providing computing time through the John von Neumann Institute for Computing (NIC) on the GCS.

References

1. E. Pavarini, Riv. Nuovo Cimento, Springer-Nature **44**, 597, 2021.
2. W. Metzner and D. Vollhardt, Phys. Rev. Lett. **62**, 324, 1989.
3. A. Georges, G. Kotliar, W. Krauth, M. J. Rozenberg, Rev. Mod. Phys. **68**, 12, 1996.
4. M. A. Korotin, A. O. Anokhin, V. I. Anisimov, A. I. Poteryaev, and G. Kotliar, J. Phys. Cond. Mat. **9**, 7539, 1998.
5. A. I. Lichtenstein and M. I. Katsnelson, Phys. Rev. B **57**, 6884, 1998.
6. E. Pavarini, E. Koch, D. Vollhardt, and A. Lichtenstein (eds.), *DMFT at 25: Infinite Dimensions*, Modeling and Simulation, Vol. 4, Verlag des Forschungszentrum Jülich, 2014, ISBN 978-3-89336-953-9.
7. E. Pavarini, E. Koch, A. Lichtenstein, and D. Vollhardt (eds.), *DMFT: From Infinite Dimensions to Real Materials*, Modeling and Simulation, Vol. 8, Verlag des Forschungszentrum Jülich, 2018, ISBN 978-3-95806-313-6.
8. A. N. Rubtsov, V. V. Savkin, and A. I. Lichtenstein, Phys. Rev. B **72**, 035122, 2005.
9. P. Werner and A. J. Millis, Phys. Rev. B **74**, 155107,2006.
10. E. Gull, P. Werner, A. J. Millis, and M. Troyer, Phys. Rev. B **76**, 235123, 2007.
11. E. Gorelov, M. Karolak, T. O. Wehling, F. Lechermann, A. I. Lichtenstein, and E. Pavarini, Phys. Rev. Lett. **104**, 22641, 2010.
12. A. Flesch, E. Gorelov, E. Koch, E. Pavarini, Phys. Rev. B **87**, 195141, 2013.
13. G. Zhang and E. Pavarini, Phys. Rev. B **104**, 125116, 2021.
14. J. Musshoff, A. Kiani, and E. Pavarini, Phys. Rev. B **103**, 075136, 2021.
15. X-J. Zhang, E. Koch and E. Pavarini, Phys. Rev. B **102**, 035113, 2020.
16. G. Zhang and E. Pavarini, Phys. Rev. B **101**, 205128, 2020.
17. G. Zhang and E. Pavarini, Phys. Rev. B **99,** 125102, 2019.
18. J. Musshoff, G. Zhang, E. Koch, E. Pavarini, Phys. Rev. B **100**, 045116, 2019.
19. E. Pavarini and E. Koch (eds.), *Simulating Correlations on Computers*, Modeling and Simulation, Vol. 11, Verlag des Forschungszentrum Jülich, 2021, ISBN 978-3-95806-529-1.
20. E. Pavarini and E. Koch (eds.), *Topology, Entanglement, and Strong Correlations*, Modeling and Simulation, Vol. 10, Verlag des Forschungszentrum Jülich, 2020, ISBN 978-3-95806-466-9.