

RESEARCH ARTICLE | AUGUST 07 2024

# KeldyshQFT: A C++ codebase for real-frequency multiloop functional renormalization group and parquet computations of the single-impurity Anderson model

Special Collection: [Algorithms and Software for Open Quantum System Dynamics](#)

Nepomuk Ritz  ; Anxiang Ge ; Elias Walter ; Santiago Aguirre; Jan von Delft ; Fabian B. Kugler 



*J. Chem. Phys.* 161, 054118 (2024)

<https://doi.org/10.1063/5.0221340>



 Nanotechnology & Materials Science  Optics & Photonics  Impedance Analysis  Scanning Probe Microscopy  Sensors  Failure Analysis & Semiconductors

Unlock the Full Spectrum.  
From DC to 8.5 GHz.  
Your Application. Measured.

[Find out more](#)



# KeldyshQFT: A C++ codebase for real-frequency multiloop functional renormalization group and parquet computations of the single-impurity Anderson model

Cite as: J. Chem. Phys. 161, 054118 (2024); doi: 10.1063/5.0221340

Submitted: 31 May 2024 • Accepted: 3 July 2024 •

Published Online: 7 August 2024



View Online



Export Citation



CrossMark

Nepomuk Ritz,<sup>1,a)</sup> Anxiang Ge,<sup>1</sup> Elias Walter,<sup>1</sup> Santiago Aguirre,<sup>1</sup> Jan von Delft,<sup>1</sup>   
and Fabian B. Kugler<sup>2</sup>

## AFFILIATIONS

<sup>1</sup>Arnold Sommerfeld Center for Theoretical Physics, Center for NanoScience, and Munich Center for Quantum Science and Technology, Ludwig-Maximilians-Universität München, 80333 Munich, Germany

<sup>2</sup>Center for Computational Quantum Physics, Flatiron Institute, 162 5th Avenue, New York, New York 10010, USA

**Note:** This paper is part of the JCP Special Topic on Algorithms and Software for Open Quantum System Dynamics.

<sup>a)</sup>Author to whom correspondence should be addressed: [nepomuk.ritz@physik.uni-muenchen.de](mailto:nepomuk.ritz@physik.uni-muenchen.de)

## ABSTRACT

We provide a detailed exposition of our computational framework designed for the accurate calculation of real-frequency dynamical correlation functions of the single-impurity Anderson model in the regime of weak to intermediate coupling. Using quantum field theory within the Keldysh formalism to directly access the self-energy and dynamical susceptibilities in real frequencies, as detailed in our recent publication [Ge *et al.*, Phys. Rev. B **109**, 115128 (2024)], the primary computational challenge is the full three-dimensional real-frequency dependence of the four-point vertex. Our codebase provides a fully MPI+OpenMP parallelized implementation of the functional renormalization group (fRG) and the self-consistent parquet equations within the parquet approximation. It leverages vectorization to handle the additional complexity imposed by the Keldysh formalism, using optimized data structures and highly performant integration routines. Going beyond the results shown in the previous publication, the code includes functionality to perform fRG calculations in the multiloop framework, up to arbitrary loop order, including self-consistent self-energy iterations. Moreover, implementations of various regulators, such as hybridization, interaction, frequency, and temperature, are supplied.

© 2024 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/5.0221340>

## I. INTRODUCTION

In the study of strongly correlated electrons, dynamical correlation functions are quantities of major interest, as they provide insights into the collective behavior and emergent phenomena arising from electronic interactions. Capturing the effects of two-particle (or four-point) correlations is one of the current major frontiers in the field. Their dynamical properties are inherently difficult to compute, as they involve three independent frequency arguments.

While most previous works on this subject focused on four-point functions in imaginary frequencies in the Matsubara

formalism<sup>2,3</sup> (MaF), obtaining real-frequency information is crucial for direct comparisons to experiments. The extraction of real-frequency data from the results of a calculation in the MaF is, in principle, possible via analytic continuation.<sup>4</sup> However, it is hard to do so reliably in practice, as the conditions for the procedure outlined in Ref. 4 are not met by finite amounts of numerical data. This renders analytic continuation an ill-defined problem, despite numerous attempts.<sup>5–7</sup> Furthermore, it had not been worked out in full detail until very recently<sup>8</sup> how analytic continuation of four-point functions could be achieved even under the assumption of analytically available results.

Pioneering attempts to directly compute real-frequency dynamical four-point correlation functions using simplified approaches made use of diagrammatic ladder approximations<sup>9,10</sup> or were restricted to a simplified frequency dependence.<sup>11–13</sup> The first fully unbiased treatment of the fluctuations contributing to the four-point vertex was achieved only a few years ago using a multipoint extension of the numerical renormalization group (NRG).<sup>14,15</sup>

Even more recently, we presented a similarly unbiased treatment of the four-point vertex of the single-impurity Anderson model using a QFT framework within the Keldysh formalism (KF), employing two related diagrammatic methods: the functional renormalization group (fRG) and the self-consistent parquet equations in the parquet approximation.<sup>1</sup> While we focused on the conceptual aspects and discussed the performance of the methods in great detail in the previous publication, here we wish to provide a detailed exposition of the computational framework for the numerical calculations of self-energies and vertex functions. In addition to what was shown in Ref. 1, the code discussed in this paper is capable of performing fRG calculations in the multiloop framework up to an arbitrary loop order, which connects the fRG to the parquet formalism.<sup>16–18</sup>

This paper aims to serve as a reference for future extensions or revisions of the code. The codebase discussed here was developed by several people over the course of multiple years, during which some goals and priorities changed and the code had to be adapted accordingly. This paper will document how the code works and what was learned during its development.

Some general design choices made during development resulted in convenient features of the code and are recommended for future projects. In the following, we briefly discuss the most important features:

a. *Modularity.* Every main building block of the code and each functionality is implemented individually, using classes and functions that serve one purpose only. As a consequence, a developer can keep an overview of the functionality. It is also comparatively easy to reuse existing features and combine them into new functionality. For example, for both the computation of the Schwinger–Dyson equation during parquet computations and the evaluation of the flow equation for the self-energy during the solution of an mfRG flow, the same classes for vertices, propagators, self-energies, and the same function for contracting a loop are used, as described in Secs. II C and II D. In addition, modularity enables unit-testing of each functionality, something too often ignored during research software development. Modularity is probably the most important feature that should be prioritized in developing any research software.

b. *Flexibility.* A modular design makes the code flexible, too. Some additional choices were made to improve its flexibility even further. Most importantly, the code enables computations in three different formalisms: the finite-temperature Matsubara formalism (MaF), the zero-temperature Matsubara formalism, and the Keldysh formalism (KF), which works at any temperature and generalizes to systems out of thermal equilibrium. Consequently, some functionality had to be implemented multiple times, such as contractions, which require summations over discrete Matsubara frequencies in the finite-temperature MaF but integrations over continuous frequencies in the zero-temperature MaF and the KF. Additionally,

in the KF, all quantities are complex-valued, whereas they are real-valued in the MaF for particle-hole symmetry. Template parameters were introduced to enable the same functions to work with objects of different types. Despite the resulting additional complexity, this conveniently enables computations in each of these three formalisms in the same codebase, still using much of the same functionality.

c. *Performance.* Computing dynamical correlation functions is a computationally demanding task, especially for four-point functions that depend on three frequency arguments. Depending on the desired resolution, this requires both excessive memory to store these functions during computations and central processing unit (CPU) power to perform computations for each combination of arguments. Concerning the latter, using optimized data structures for efficient readouts of data as well as an efficient but still precise algorithm for integrating over frequencies (the numerical bottleneck) improved matters significantly. In addition, using a compiled programming language is basically a must, and keeping track of constant variables and member functions helps the compiler optimize the code.

d. *Scalability.* Apart from the simplest calculations, most diagrammatic calculations would not be feasible without parallelization. This is because practically all calculations in parquet formalism, or mfRG, require computations for all possible combinations of external arguments of the correlation functions. As those are independent from each other, it is possible and advisable to parallelize the demanding computations of bubbles and loops (see Sec. II D) in the external arguments. Using the OpenMP and MPI interfaces, this can easily be achieved for parallelization across different threads on the same node and across multiple nodes, respectively (for more details, see Sec. II G 1). As long as the memory requirements are met, the performance of the code scales almost perfectly with the computational resources.

At this point, we disclose that the present code also has a number of weaknesses that evolved over the course of development. If the reader intends to set up a new codebase for the purpose discussed here, we recommend considering the following points:

a. *Too many preprocessor macros (“flags”).* The code contains far too many preprocessor macros, used to specify different parameters and settings before compilation (see Sec. II I). This not only hampers readability but also increases the risk of errors, as it is never possible to test the full functionality of the code because one would have to compile and test all possible configurations independently. With simple combinatorics, this quickly becomes an overwhelming task. Using preprocessor macros is, however, useful for quick implementations of new functionality, which is why they accumulate over time.

b. *Too many overly complicated structures.* The code contains several classes that are way more complicated than they need to be, such as the different vertex classes or the data buffer (see Secs. II C 1 and II G 8). When they were set up, the goal was to keep them as general as possible, such that they could be used for all kinds of models in all kinds of formalisms. For this purpose, templates are used excessively as well. As a consequence, they are indeed flexible, but they are cumbersome to use in any specific context, and their implementations are difficult to grasp. In addition, the code takes a long time to compile and link, which is inconvenient for everyday

development. Ultimately, as a developer, one has to find the right trade-off between flexibility and simplicity.

c. *Too little use of existing implementations.* Several textbook algorithms, such as the Gauß-Lobatto routine for frequency integrations or the Cash-Karp routine for solving ODEs (see Secs. II G 5 and III C 1), were implemented by hand. The reason for this was the desire to comprehend and track the inner workings of the algorithms at every point during a calculation. In hindsight, much time and effort could have been saved if existing implementations of these algorithms had been used as “black boxes.”

d. *Language.* C++ is a very versatile language that runs on essentially any computer and can produce very fast code. However, a codebase written in C++ requires a lot of work to write and maintain. Initially, C++ was chosen for performance reasons. By now, however, there are established alternative programming languages that are easier to use, less error-prone, and (almost) as fast, such as Julia,<sup>19</sup> Rust,<sup>20</sup> or Mojo.<sup>21</sup>

e. *Priorities.* Driven by the desire to obtain data with maximal resolution and precision, the top priority has always been performance. While this is very typical for codes written by physicists, it is not in line with the typical recommendation in software engineering, which would prioritize correctness and maintainability *over* performance.<sup>22</sup> While we are confident that the code produces correct results after extensive benchmarks,<sup>1</sup> the code is not written in the simplest way and is not easily readable and maintainable. While we acknowledge that generating results quickly is deemed to be the most important aspect of research at present, we advocate for reconsidering the priorities during research software development for future projects.

The rest of the paper is structured as follows: In Sec. I A, we briefly introduce the single-impurity Anderson model (AM). In Sec. I B, we briefly recapitulate the main concepts of diagrammatic many-body theory. In Sec. I C, we comment on the complications that arise by performing computations in the very general Keldysh formalism, which is the main selling point of the present codebase.

In the second part of the paper, we give details on the code itself, introducing the main objects in Sec. II C and explaining the main functionality in Sec. II D. We list several options for postprocessing the raw data obtained after a completed calculation in Sec. II E and briefly explain how the data are organized in Sec. II F. Special emphasis is placed on performance-critical aspects of the code in Sec. II G. We comment on how the code is tested in Sec. II H. Finally, we provide an overview of the most important options for parameter choices that can be performed in Sec. II I, illustrating the versatility of the codebase.

In the third main part of the paper, we elaborate on how three different diagrammatic algorithms, perturbation theory, the parquet equations, and the mfRG, are implemented. In particular, we list the different flow schemes that are available in mfRG. Finally, Sec. IV presents a conclusion.

Before the end of this introduction, a disclaimer is in order: This paper does not mention every single class or function in the code but focuses on the most important aspects and functionalities. In addition, while the code enables computations in the KF and the MaF at both finite and zero temperatures, we focus our specific descriptions mainly on the KF functionality, as this is a unique feature of our codebase.

## A. Model

We consider the single-impurity Anderson model (AM) in thermal equilibrium, one of the most studied models in all of condensed matter physics. Its physical behavior is well understood, and numerically exact benchmark data for single-particle correlation functions is available from NRG,<sup>23</sup> as are exact analytical results for static quantities at zero temperature from the Bethe ansatz.<sup>24,25</sup> This makes it an ideal candidate for studies focused on reliable method development.

The AM is a minimal model for localized magnetic impurities in metals introduced by Anderson to explain the physics behind the Kondo effect.<sup>26</sup> It is defined by the Hamiltonian

$$H = \sum_{\epsilon\sigma} \epsilon c_{\epsilon\sigma}^\dagger c_{\epsilon\sigma} + \sum_{\sigma} \epsilon_d n_{\sigma} + U n_{\uparrow} n_{\downarrow} + \sum_{\sigma} (V_{\epsilon} d_{\sigma}^\dagger c_{\epsilon\sigma} + \text{H.c.}), \quad (1)$$

describing a local impurity  $d$  level with on-site energy  $\epsilon_d$ , hybridized with spinful conduction electrons, created by  $c_{\epsilon\sigma}^\dagger$ , of the metal via a matrix element  $V_{\epsilon}$ . Hence, it qualifies as an open quantum system. The electrons in the localized  $d$  state, where  $n_{\sigma} = d_{\sigma}^\dagger d_{\sigma}$ , interact according to the interaction strength  $U$ , whereas the  $c$  electrons of the bath are non-interacting. The bath electrons are hence formally integrated out, yielding the frequency-dependent retarded hybridization function  $-\text{Im} \Delta^R(\nu) = \sum_{\epsilon} \pi |V_{\epsilon}|^2 \delta(\nu - \epsilon)$ . We consider a flat hybridization in the wideband limit,  $\Delta^R(\nu) = -i\Delta$ , so that the bare impurity propagator reads  $G_0^R(\nu) = (\nu - \epsilon_d + i\Delta)^{-1}$ .

The code can treat all choices for the on-site energy  $\epsilon_d$ . For the special choice  $\epsilon_d = -U/2$ , the model has particle-hole symmetry and is referred to as the symmetric Anderson model (sAM). This setting simplifies the calculations somewhat. For instance, in this case, the Hartree-term of the self-energy is constant  $\Sigma_H = U/2$  (see also Sec. III A 1). In addition, in the MaF, all quantities become real-valued, whereas they are complex-valued otherwise. Hence, the code supplies a parameter flag to make use of these properties (see Sec. II I). For general  $\epsilon_d \neq -U/2$ , we speak of the asymmetric Anderson model (aAM).

Some physical applications require an additional external magnetic field  $h$ , described by an additional term  $h(n_{\uparrow} - n_{\downarrow})$  in the Hamiltonian. At present, the codebase is, however, not applicable in this setting, as this would break  $SU(2)$  symmetry, which is heavily used and hard-coded into the codebase (see Sec. II G 3). A generalization to  $h \neq 0$  is possible but would require major effort.

While the present implementation is restricted to the AM, the code in principle can also treat other models: all data structures possess an additional internal index suitable for encoding additional dependencies and quantum numbers of more complicated models, such as a momentum dependence or multiple orbitals. Indeed, the first attempts to study the 2D Hubbard model had been started; however, the simplest KF perturbation theory calculations turned out to be too demanding at the time. The corresponding functionality is, therefore, not included in this release.

## B. Diagrammatic many-body theory

The basic objects of interest in all our calculations are one- and two-particle correlation functions. Their non-trivial contributions due to interaction effects are contained in the self-energy  $\Sigma$  and the four-point vertex  $\Gamma$ ,

$$\Sigma = \text{---} \circ \Sigma \text{---}, \quad \Gamma = \text{---} \square \Gamma \text{---} \quad (2)$$

The self-energy is used together with the bare propagator  $G_0$  to express the one-particle propagator  $G$  via the Dyson equation

$$G = \text{---} = \text{---} + \text{---} \circ \Sigma \text{---}, \quad (3)$$

which is formally solved by  $G = 1/(G_0^{-1} - \Sigma)$ . The vertex is the connected and amputated part of the two-particle correlation function  $G^{(4)}$ ,

$$G^{(4)} = \text{---} \text{---} - \text{---} \text{---} + \text{---} \square \Gamma \text{---} \text{---}, \quad (4)$$

from which physical susceptibilities can be obtained by contracting pairs of external legs (see Appendix C of Ref. 1 for details). The first-order contribution to the vertex is given by the fully antisymmetric, local, and instantaneous bare vertex, represented as a single dot,

$$\Gamma_0 = \text{---} \times \text{---} \sim U > 0, \quad (5)$$

in standard Hugenholtz notation. Using the bare vertex and the bare propagator  $G_0$ , diagrammatic perturbation series for both the self-energy and the vertex can be derived, which will be the subject of Sec. III A. A perturbation series up to finite order in  $\Gamma_0$  is, however, only appropriate for weak coupling strengths. In order to reach larger couplings, an infinite number of diagrams have to be summed. This is the purpose of two related formalisms, the parquet formalism and the multiloop functional renormalization group, to be discussed in Secs. III B and III C, respectively. Both formalisms employ the parquet decomposition to organize all diagrammatic contributions to  $\Gamma$  into one of four distinct categories: Two-particle reducible diagrams in one of the three two-particle channels  $a$ ,  $p$ , and  $t$ , included in the three two-particle reducible vertices  $\gamma_{r \in \{a,p,t\}}$  or two-particle irreducible diagrams, included in the fully irreducible vertex  $R$ ,

$$\Gamma = \text{---} \square R \text{---} + \text{---} \square \gamma_a \text{---} + \text{---} \square \gamma_p \text{---} + \text{---} \square \gamma_t \text{---} \quad (6a)$$

$$= \text{---} \times \text{---} + \text{---} \text{---} + \frac{1}{2} \text{---} \text{---} - \text{---} \text{---} + O[(\Gamma_0)^3]. \quad (6b)$$

Any specific diagram is said to be two-particle reducible if it can be disconnected by splitting a propagator pair. Otherwise, it is said to be two-particle irreducible. The parquet decomposition is exact, as it in essence just provides a classification of all diagrams that contribute to  $\Gamma$ . However, neither the parquet formalism nor the mFRG provide equations for  $R$ . In practice, some approximation is required. The simplest one is the parquet approximation (PA)

$$R = \Gamma_0 + O[(\Gamma_0)^4] \approx \Gamma_0, \quad (7)$$

which approximates the fully irreducible vertex  $R$  by the bare vertex  $\Gamma_0$ . As it introduces an error in the fourth order in perturbation theory, it fails for large coupling strengths and is hence applicable only up to intermediate couplings. The PA was applied throughout in Ref. 1 and is the only one so far implemented in the codebase (see Sec. II C 1 for a comment on other possibilities).

### C. Keldysh formalism

The following section assumes familiarity with the KF and describes challenges arising for computations with the KF rather than the more widespread MaF (for a more extensive discussion of the KF, see Refs. 27 and 28).

The KF<sup>29–31</sup> works both out of equilibrium and in thermal equilibrium at arbitrary temperature, in a real-frequency description. This is an advantage over the more popular MaF, which works at imaginary (“Matsubara”) frequencies, requiring analytical continuation, a mathematically ill-defined problem if one works with a finite amount of imperfect numerical data. Still, the KF is seldomly used because practical calculations are more complicated for two main reasons.

In the KF, all operators acquire an additional contour index, which specifies whether they sit on the forward or backward branch of the Keldysh double-time contour. It follows that the four-point vertex, for example, has  $2^4 = 16$  different components. While some of these components can be eliminated by causality or related to other components by fluctuation–dissipation relations in thermal equilibrium or symmetries, this additional index structure complicates the implementation and the numerics.

In thermal equilibrium, energy conservation can be leveraged by Fourier-transforming all correlation functions into frequency space. In contrast to the MaF at finite temperatures, this dependence is continuous. Hence, contractions over frequency arguments require numerically more expensive integrations instead of summations. The integrations become more costly at lower temperatures as the frequency dependence of the correlation functions becomes more sharply peaked. The four-point functions, which depend on three continuous frequency arguments, are the numerical bottleneck for which arbitrarily high resolutions are out of reach due to both computation and memory demands. Discretizing the frequency dependence in a clever way and using adaptive integration routines is, therefore, key, as discussed in Secs. II G 4 and II G 5.

Finally, the KF also allows for computations outside of thermal equilibrium. However, the present discussion is restricted to thermal equilibrium. Extending the code out of equilibrium is possible with moderate effort.

## II. THE CODE

In part II of the paper, we describe the main building blocks of the code—the classes representing correlation functions and other functions for combining them in diagrammatic computations. Furthermore, we describe post-processing schemes and emphasize aspects important for performance. More information on the technical details of individual code pieces can be found in the documentation attached to the source code (see the code availability statement at the end of this paper).

## A. Prerequisites

The code itself is written in C++<sup>17,32</sup> and is built using CMake,<sup>33</sup> demanding at least version 3.10. It requires the GSL,<sup>34</sup> boost,<sup>35</sup> and Eigen3<sup>36</sup> libraries, as well as the HDF5<sup>37</sup> library for input and output. For parallelization, the OpenMP<sup>38</sup> and MPI<sup>39</sup> interfaces are used. Notably, we do not supply precompiled executables that could be run directly, for several reasons: First, the code makes heavy use of preprocessor flags that must be set *before* compilation and that are in part used to specify the concrete problem at hand (see Sec. II I). Second, special compilers for the particular architecture at hand might be available, which could optimize the code during compilation and linking, improving the performance. The user should hence adapt the file CMakeLists.txt accordingly, such that the required libraries are included and linked properly and all compiler settings are as desired.

The technical documentation supplied with the code is generated automatically using the tools Doxygen,<sup>40</sup> Sphinx,<sup>41</sup> Breathe,<sup>42</sup> and CMake.

## B. Basic structure

The structure of the main part of the codebase is depicted in Fig. 1. The main objects of interest are the SelfEnergy  $\Sigma$  and the four-point Vertex  $\Gamma$ . Separate classes have been implemented for both, discussed in detail below. Both classes use instances of the class that defines suitably chosen FrequencyGrids, to be discussed in Sec. II G 4, for discretizing the continuous frequency dependence. A self-energy and a vertex always come together in any practical calculation, representing data for a step of an mFRG flow or an iteration of the parquet solver. The self-energy and vertex classes are hence combined in a State class  $\Psi = (\Sigma, \Gamma)$ . The algorithms discussed in Sec. III require computing bubble- and loop-type diagrams, the main functionality of the codebase. As detailed in Sec. II D 1 below, the bubble\_function contracts two input vertices with a pair of propagators in one of the three two-particle channels to yield a new four-point vertex, which is stored as an instance of the Vertex class. For example, contracting two vertices  $\Gamma_1$  and  $\Gamma_2$  in the *a* channel is denoted as

$$B_a(\Gamma_1, \Gamma_2) = \Gamma_1 \circ \Pi_a \circ \Gamma_2 = \text{[Diagram of bubble diagram]}, \quad (8)$$

see also Appendix C in Ref. 1 for a fully parametrized version. The required propagator pair  $\Pi$  belongs to a separate Bubble

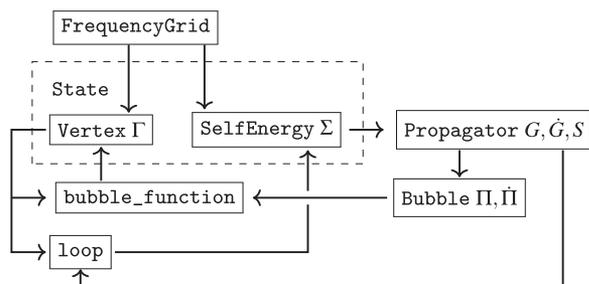


FIG. 1. Schematic depiction of the main parts of the codebase.

class, ensuring the correct combination of propagators and their parametrization. The propagators themselves are defined in the Propagator class, which essentially implements the Dyson equation, Eq. (3), combining  $G_0$  and  $\Sigma$ . The former contains all the system parameters, including the regulator in mFRG; the latter encodes the interaction effects. Both the Propagator and Bubble classes can handle differentiated objects arising in mFRG (see Sec. III C). Finally, the loop function is used to contract two external legs of a four-point Vertex with a Propagator, yielding an instance of the SelfEnergy class, for example,

$$L(\Gamma, G) = \text{[Diagram of loop diagram]} \quad (9)$$

These types of diagrams are required, e.g., for the mFRG flow equation of the self-energy or for the evaluation of the SDE after a previous bubble diagram computation.

## C. Correlation function classes

In the following, we discuss the main building blocks of the code in more detail. We begin by outlining the self-energy and vertex classes. In addition, there are two helper classes: the first represents propagators, combining the bare propagator and the self-energy; the second combines a pair of propagators as needed for bubble-type diagrams.

### 1. The Vertex classes

In total, the code contains the four classes irreducible, rvert, fullvert, and GeneralVertex to store different types of four-point vertices.

The irreducible class contains the two-particle irreducible part of the vertex,  $R$ . In the PA, its 16 Keldysh components are just constants. It can easily be extended to hold nontrivial input data, for example, in the context of diagrammatic extensions<sup>43</sup> of dynamical mean-field theory<sup>44</sup> such as DFA<sup>45</sup> or DMF<sup>2</sup>RG.<sup>46</sup>

The rvert class stores the two-particle reducible vertices  $\gamma_{re \{a,p,t\}}$ . Each of them is split up into their asymptotic classes,<sup>47</sup>  $K_1$ ,  $K_2$ , and  $K_3$ , where the  $K_{2'}$  class is inferred from  $K_2$  by crossing symmetry. Being one-, two-, and three-dimensional objects, respectively, each of those naturally has its own frequency grid. The rvert supplies several methods to store and read out data, either directly or interpolated. Conveniently, it can return all vertex parts where external legs either do or do not meet at the same bare vertex on the left or on the right-hand side by suitably combining the  $K_1$ ,  $K_{2'}$ , or  $K_{2''}$  and  $K_3$  classes, respectively. This turned out to be very handy for keeping track of contributions for the different asymptotic classes during calculations. In addition, the rvert class can track and, if desired, enforce symmetries in the Keldysh-, spin-, and frequency domains (see Sec. II G 3 for details). For debugging purposes, functionality not using symmetries is provided as well.

The fullvertex class combines one instance of the irreducible class and three instances of the rvert class, one for each two-particle channel  $a, p, t$ . It can then return the value of the full vertex, which is the sum of the four contributions for a given

Keldysh and spin component, interpolated at a given combination of frequencies. As each individual `rvert` instance, it can collect all those parts of the vertex where the external legs either do or do not meet at the same bare vertex on the left or on the right-hand side and includes functionality to exploit various symmetries. In addition, it can compute the  $p$ -norm of each asymptotic contribution, which is useful for debugging purposes and convergence criteria, e.g., in parquet computations.

While instances of the `fullvertex` class hold the data of the symmetry-reduced sector of a full vertex, certain diagrammatic equations involve subsets of vertex diagrams. One example is the  $r$ -channel-irreducible vertex used in the Bethe–Salpeter equations outlined in Sec. III B. Such diagrams do not necessarily obey all the symmetries of a full vertex, so they must be treated differently. These asymmetric cases are, therefore, encoded in the `GeneralVertex` class. It uses multiple instances of `fullvertex`, which together cover the symmetry-reduced sector of the asymmetric vertex data. Let us comment here that, while this approach is feasible, it turned out to be inconvenient in practice, as one always has to make sure that all sectors are covered, i.e., that all required `fullvertex` instances are provided. This is a source of logical errors that can sometimes be hard to find. In retrospect, it would have been better to pay the increased cost in memory to store all vertex contributions in the same object, making the code easier to read and to work with.

All vertex classes allow adding or subtracting two instances of the respective classes or multiplying a number with a vertex instance.

Splitting up the vertex functionality into so many different classes was made at the beginning of developing the code to provide enough flexibility, in particular regarding symmetries and a possible non-trivial input for the irreducible vertex. In hindsight, it turned out that for the computations performed in Ref. 1, this structure would not have been required in this generality.

## 2. The `SelfEnergy` class

The `SelfEnergy` class comes with a `dataBuffer` that stores the discrete values of the retarded and Keldysh components of the self-energy on a given frequency grid (see Secs. II G 4 and II G 8). When instantiating an object of the `SelfEnergy` type, a given frequency grid can either be supplied or a suitable one is generated automatically based on the value of the regulator  $\Lambda$ . In addition, the asymptotic value of the retarded component of the self-energy has to be set. Most of the time, this should be the Hartree value  $\Sigma_H$ , as the `SelfEnergy` inside the code is supposed to be used only for the dynamical, i.e., frequency-dependent, contributions of the self-energy, which excludes the constant Hartree value. For the sAM, the Hartree value is constant,  $\Sigma_H = U/2$ ; in the asymmetric case, it has to be computed self-consistently beforehand. This can be performed inside the code using the `HartreeSolver` class (see Sec. III A 1).

The `SelfEnergy` class provides a host of methods used throughout the code. Most importantly, it can return the value of the self-energy either directly at a given input on the frequency grid (fast) or return an interpolated value at a given continuous frequency (not so fast). It can also set the value of  $\Sigma$  for a given input. In addition, one can compute the  $p$ -norm of  $\Sigma$  and the relative deviation to a different `SelfEnergy` instance using the maximum norm. This

is used to check convergence in parquet computations detailed in Sec. III B.

Finally, multiple operators are defined for the `SelfEnergy` class, which are used to add or subtract two `SelfEnergy` instances or to multiply some number with a `SelfEnergy` instance.

## 3. The `State` class

Instances of the `State` class are the high-level objects that are mainly used by the high-level algorithms discussed in Sec. III. The `State` class combines a `GeneralVertex` and a `SelfEnergy`, which together contain all non-trivial information that one might wish to compute. In that sense, it suffices to completely specify the “state” of the calculations. For the purpose of fRG calculations, the `State` class also holds the value of the flow parameter  $\Lambda$ .

As with the vertex classes and the `SelfEnergy` class, the `State` class also comes with operators that can be used to add and subtract states from one another and to multiply a number with a state. Under the hood, these operators just invoke the corresponding operators previously defined for the vertex and self-energy. Hence, all high-level algorithms can manipulate instances of the `State` class directly, e.g., by combining several iterations of the parquet solver in a mixing scheme.

## 4. The `Propagator` class

The `Propagator` class is special in the sense that it stores almost no data itself. Instead, it references instances of the `SelfEnergy` class and combines the analytical form of the bare propagator  $G_0$  with the self-energy via the Dyson equation,  $G = 1/[(G_0)^{-1} - \Sigma]$ . To that end, it can return the value of a given propagator at some point, interpolated on the frequency grid of the referenced self-energy. This can be performed either directly for a given Keldysh component at some continuous frequency or vectorized over all Keldysh components. As  $G_0$  depends on the formalism used and in mfRG on the choice of the regulator, separate methods for a variety of choices are provided. In addition, one can specify whether the full propagator  $G$  shall be computed, or the single-scale propagator  $S$ , the differentiated propagator including the Katanin extension,<sup>48</sup> or just the Katanin extension by itself (see Sec. III C). Note that the Katanin extension requires the self-energy differentiated with respect to the flow parameter  $\Lambda$ ; hence, the propagator class references *two* `SelfEnergy` instances, one non-differentiated and one differentiated.

## 5. The `Bubble` class

Finally, the `Bubble` class combines two propagators to yield a bubble in one of the three two-particle channels  $a$ ,  $p$ , and  $t$ , according to Eqs. (C1a)–(C1c) in Ref. 1. For evaluating differentiated bubbles in mfRG, one of the propagators can be chosen to be the single-scale propagator  $S$  or the fully differentiated one  $\hat{G}$ . In that case, the bubble already takes care of the product rule, giving (symbolically)  $\hat{\Pi}^S = GS + SG$  or  $\hat{\Pi} = G\hat{G} + \hat{G}G$ . Otherwise, it just yields  $\Pi = GG$ . The `Bubble` class provides functions for obtaining the value of a bubble in a given channel at specified bosonic and fermionic frequencies, either for one specific Keldysh component directly or vectorized over the Keldysh structure. This class simplifies bubble computations using the `bubble_function` (see Sec. II D 1).

## D. Main functions for diagrammatic computations

Computing bubbles and loops involves contractions over quantum numbers and Keldysh indices, including integrations over frequencies for all possible combinations of external arguments, and is by far the most costly part for the numerics. A clean and efficient implementation of this functionality is, therefore, paramount and should be of the highest priority when setting up a new code. In the following, we provide technical details on this most important part of the code.

### 1. The bubble\_function

The `bubble_function` implements Eqs. (C2a)–(C2c) from Ref. 1. It takes references to three vertices as arguments, one to store the result of the computation and two others to be connected by a `Bubble` object. This `Bubble` object can either be supplied as well or is initialized by an overload of the `bubble_function`, which in addition requires the two propagators that shall be used for the `Bubble`. The main work is then performed by an instance of the class `BubbleFunctionCalculator`, which performs the bubble contractions for each diagrammatic class separately. This is performed for every possible combination of external arguments, i.e., Keldysh indices and frequencies. At this point, the calculations are parallelized as outlined in Sec. II G 1. For each set of arguments, an `Integrand` object is instantiated, which puts together the two vertices and the bubble and performs the contraction over Keldysh indices if the flag `SWITCH_SUM_N_INTEGRAL` is set to 1. The `Integrand` class provides an operator that reads out the integrand at a given frequency. It is called by the integrator, invoked subsequently, and described in detail in Sec. II G 5. The results of all the frequency integrations are finally collected and *added* to the vertex object that was given as the first argument to the `bubble_function`. The choice not to output a completely new vertex but instead to add the result to an existing vertex has historical reasons to save memory. This increased the risk of logical errors during high-level algorithm implementations, though, and in hindsight, the `bubble_function` should better have been designed to output a completely new vertex object.

### 2. The loop function

The `loop function` implements Eq. (C3) from Ref. 1 and is structured similarly to the `bubble_function`. It takes a reference to self-energy for storing the result as well as references to a vertex and a propagator as arguments for the loop. For each external fermionic frequency, in which the computation is parallelized again, it invokes the integrator to perform a frequency integration using the `IntegrandSE` class. For the aAM, the asymptotic value of the just computed self-energy is extracted from the Hartree- and the  $K_{1,t}$  and  $K_{2',t}$  terms after the calculation. For the sAM, the asymptotic value of the self-energy is a known constant.

## E. Postprocessing

The code provides a host of postprocessing functions. These are not required for the actual calculations themselves but are useful to extract additional information from their results, either as consistency checks or to infer derived quantities for later analysis.

### 1. Causality check for the self-energy

By causality, the imaginary part of the retarded component of the self-energy is strictly non-positive;<sup>49</sup>  $\text{Im} \Sigma^R(\nu) \leq 0$  for all frequencies  $\nu \in \mathbb{R}$ . A violation of this condition not only constitutes an unphysical result but often leads to numerical instabilities. The code, therefore, provides the function `check_SE_causality` that checks this condition for a supplied instance of `SelfEnergy`. Typically, this function is invoked after each ODE step during an mFRG calculation or after each iteration of the parquet solver.

### 2. Fluctuation dissipation relations

In thermal equilibrium at temperature  $T$ , one has a fluctuation–dissipation relation (FDR)<sup>11,27</sup> between the retarded and the Keldysh components of the propagator,  $G^K(\nu) = 2i \tanh(\frac{\nu}{2T}) \text{Im} G^R(\nu)$ , and the self-energy,  $\Sigma^K(\nu) = 2i \tanh(\frac{\nu}{2T}) \text{Im} \Sigma^R(\nu)$ . This relation can be used to infer the Keldysh components of the self-energy from the retarded component or vice versa; hence, it would in principle suffice to compute only one of the components. However, in the vectorized form of the code, both components of the self-energy are computed anyway. The FDR can hence be used as an internal consistency check, provided by the function `check_FDTs_selfenergy`. It computes  $\Sigma^K$  from  $\Sigma^R$  via the FDR and compares it to the independently computed Keldysh-component of the self-energy by computing the 2-norm of the difference.

As an additional consistency check, the fulfillment of fluctuation–dissipation relations for the  $K_1$  classes, reading

$$\text{Im} K_1^R(\omega) = -\frac{i}{2} \tanh\left(\frac{\omega}{2T}\right) K_1^K(\omega), \quad (10)$$

can be examined. One may also want to check generalized FDRs for three-point and four-point contributions of the vertex.<sup>50</sup>

### 3. Kramers–Kronig relation

For functions  $f(\omega)$  that are analytic in the upper half plane, like retarded single-particle correlation functions, the Kramers–Kronig transform relates the real and imaginary parts via

$$\text{Re} f(\omega) = \frac{1}{\pi} \mathcal{P} \int_{-\infty}^{\infty} d\omega' \frac{\text{Im} f(\omega')}{\omega' - \omega}, \quad (11)$$

where  $\mathcal{P}$  denotes the Cauchy principal value. Inside the code, the function `check_Kramers_Kronig` can be used to test how well this generic analytic property is fulfilled.

### 4. Sum rule for the spectral function

The fermionic spectral function  $A(\nu) = -\text{Im} G^R(\nu)/\pi$  must obey the sum rule

$$\int_{-\infty}^{\infty} d\nu A(\nu) = 1. \quad (12)$$

The function `sum_rule_spectrum` implements this integral as a consistency check.

## 5. Susceptibilities

Susceptibilities, which are of significant physical relevance, are derived from the vertex by contracting pairs of external legs. Diagrammatically, the formula for the  $a$ -channel susceptibility reads

$$\chi_a = \text{bubble} + \text{bubble} \left[ \text{box} \right] \text{bubble}, \quad (13)$$

and similarly for the susceptibilities in the  $p$  and  $t$  channel. The fully parametrized equations are provided in Eq. (C7) of Ref. 1. Linear combinations of these diagrammatic susceptibilities yield the physical susceptibilities [see Eq. (C8) of Ref. 1]. The code computes susceptibilities using the function `compute_postprocessed_susceptibilities`, which can be invoked after a completed calculation using the name of the file that stores the results. It iterates through all layers that correspond to ODE steps or parquet iterations (see Sec. II F), evaluates Eq. (C7) using the vertex and self-energy for each, and stores the results as a new dataset in the same file.

It was found in Ref. 47 that for converged parquet computations, susceptibilities can more easily be extracted directly from the  $K_1$  class. As discussed in Refs. 1 and 51, one can also choose to compute susceptibilities that way during fRG computations, even though the two schemes are inequivalent if multiloop convergence is not reached. The two different schemes of computing susceptibilities can then be used to gauge the quality of the truncation.

## 6. Vertex slices

Finally, the function `save_slices_through_fullvertex` can be used to read out two-dimensional “slices” of the full vertex. It takes the filename corresponding to the results of a finished calculation as an argument, iterates through all layers, and saves a two-dimensional cut of all Keldysh components of the full vertex in the  $t$ -channel parametrization for zero transfer frequency ( $\omega_t = 0, v_t, v_t'$ ) for a given spin component. While this function does not perform any non-trivial calculations, it is useful for visualization purposes. If desired, the function can be straightforwardly adapted to store vertex slices at finite transfer frequencies, enabling full scans through the three-dimensional structure of the four-point vertex.

## F. I/O

We use the HDF5 file format<sup>37</sup> for input and output purposes throughout. To organize the data for output, the contents of a state are split into different datasets that correspond, e.g., to all the asymptotic classes of the vertex in each channel, the self-energy, the frequency grids, and the most important parameters of the calculation. The output file is then organized on a high level in terms of “ $\Lambda$  layers,” the idea being that each layer enables access to a different state stored in the same file. Thereby, a single file contains, e.g., the results of a full mFRG flow, where each “ $\Lambda$  layer” corresponds to a different value of the regulator. Alternatively, this structure can be used to store the results of all iterations needed for solving the parquet equations. Of course, one can equally well use just a single layer to store the end result of a computation, such as a converged solution of the parquet equations or the result of a PT2 computation.

The function `write_state_to_hdf` creates a new file with a fixed number of layers and saves an initial state into the first layer. Additional states generated during subsequent computations can be added to the same file (but into a different layer to be specified) using the function `add_state_to_hdf`. In effect, these functions are wrappers of a host of additional functions that are able to store various data structures, such as scalars, vectors, or even Eigen-matrices, in an HDF file.

When using parallelization, as detailed in Sec. II G 1, one has to ensure that only one single process writes data into the output file. Collisions, where multiple processes simultaneously try to write to the same location in memory, will cause the program to crash.

It is possible to read data from an existing HDF file to generate a new state for subsequent computations. For this purpose, the function `read_state_from_hdf` reads a state from a specified layer of a provided HDF file. One can thus do checkpointing: If all steps of an mFRG flow or all iterations of the parquet solver are stored separately, a computation that was interrupted can be continued from the last step stored. This design feature is useful for large computations that have to be split over several separate jobs or in the case of a hardware error causing a job to crash. Setting up checkpointing functionality is, therefore, strongly recommended.

## G. Performance

In the following, we discuss parts of the code of special importance for performance. Of course, there is always a trade-off between accuracy and performance, as, e.g., an arbitrary high frequency resolution quickly becomes prohibitive. Nevertheless, efficient implementations are necessary for challenging computations.<sup>52</sup> For the precision-focused calculations for which this codebase was developed, these parts are, therefore, of utmost importance.

### 1. MPI+OpenMP parallelization

As mentioned in the beginning, mFRG and parquet computations can be heavily parallelized since the correlation functions are (repeatedly) evaluated independently for every possible combination of external arguments. Parallelization is especially advisable for computing bubbles of two four-point vertices, as outlined in Sec. II D 1. We use the OpenMP interface for parallelization across multiple threads on a single node and the MPI interface for parallelization across multiple nodes. While OpenMP parallelization works with shared memory, meaning that all threads have access to the same data on the node that they are running on, one has to be careful with MPI parallelization working on distributed memory. Processes that run on different nodes to compute, say, a four-point vertex for different sets of external arguments cannot write their results into the same instance of a four-point vertex. Hence, we introduce additional buffers distributed across the nodes. After the computation of, say, a four-point vertex is finished, these buffers are collected, and their contents are put together to yield the full result. While this scheme is initially somewhat cumbersome to set up, it pays off tremendously, as the code’s performance scales well with the computational resources, including multiple nodes. This is because, first, computations for different external arguments are independent from each other, so there is minimal communication between the nodes. Second, the number of external arguments required for

precision-focused calculations is large, so individual threads have little downtime waiting for other threads to finish. For example, the most expensive calculations in Ref. 1 involved 125 points along each of the three frequency axes, which were parallelized across 32 nodes running 32 threads each. Provided enough CPU power, the resolution could, in principle, be increased further but is ultimately limited by memory.

## 2. Vectorization

As outlined in Sec. I C, KF calculations require computing  $2^n$  Keldysh components of  $n$ -point functions. These components can be arranged into a matrix, yielding, e.g., a  $4 \times 4$  matrix for the four-point vertex. This structure can be exploited for summing over Keldysh indices by using vectorization and the data structures of the Eigen library,<sup>36</sup> significantly improving performance. This works because all Keldysh components are stored in contiguous sections of memory. Of course, the other parts of the code have to be able to use these data structures properly, which is why all functions that enable, e.g., access to the correlation functions (see Sec. II C) have two versions: one that can handle matrix-valued data when vectorization is used, and another used otherwise.

When using vectorization, all Keldysh components have to be stored explicitly. As a consequence, identities that relate different Keldysh components, such as certain symmetries or FDRs, cannot be used to reduce the numerical effort. Although maximal exploitation of symmetries initially was one of our main objectives, we later found that vectorization over Keldysh components is preferential despite the larger memory costs.

In the finite- $T$  MaF, we use vectorization to represent the Matsubara frequency dependence of all correlation functions. This leads to massive speedups when performing Matsubara sums as matrix-multiplications.

## 3. Symmetries

Many symmetries for reducing the number of data points to be computed directly can still be used together with vectorization over Keldysh indices. These include crossing symmetry of the vertex, which relates a vertex to itself with one pair of external fermionic legs exchanged; complex conjugation of the vertex; SU(2) symmetry in the absence of a magnetic field (which, in combination with crossing symmetry, reduces the number of independent spin components to 1); and frequency symmetries in the presence of particle-hole symmetry. For explicit details on these symmetries, see Appendix A in Ref. 1.

Since frequency integrations are the most costly part of the computations, symmetry operations are not used for evaluating integrands on the fly. Instead, they are used to reduce the number of vertex components to be computed. Since the vectorized version of the code performs sums over Keldysh indices by matrix multiplication, the result of the integration contains all Keldysh components. Hence, we use the symmetry relations to reduce the other arguments, i.e., spin and frequency. Information about the symmetry-reduced components is encoded in symmetry tables. These contain entries for every channel, asymptotic class, spin component, and frequency sector and indicate whether a data point belongs to the symmetry-reduced sector or, otherwise, how to retrieve a value via symmetry relations.

## 4. Frequency grids

For numerical calculations, the continuous frequency dependence of correlation functions in the KF (and in the MaF at  $T = 0$ ) must be discretized. Since these functions can become sharply peaked around certain frequencies, especially at lower temperatures, but simultaneously decay only slowly asymptotically (typically  $\sim 1/\nu^2$  or even  $\sim 1/\nu$  for some components), finding a suitable discretization that resolves all sharp structures but still captures the asymptotic decay is hard. Since the sharp features mostly occur at smaller frequencies (measured relative to the hybridization  $\Delta$ ), we use a frequency grid that provides high resolution at small frequencies and fewer points at high frequencies. To achieve this, an equidistant grid of an auxiliary variable  $\Omega \in [-1, 1]$  is mapped to frequencies according to  $\nu(\Omega) = A\Omega|\Omega|/\sqrt{1 - \Omega^2}$ . The parameter  $A > 0$  can be suitably chosen automatically or by hand for all quantities, as further explained in Appendix G of Ref. 1. However, we do not recommend optimizing  $A$  automatically, as this can become expensive and unreliable in the presence of numerical artifacts.

The frequency grid is implemented in the `FrequencyGrid` class. It specifies the grid parameters such as the number of grid points or the scale factor  $A$ , and can access both continuous frequencies  $\nu$  and auxiliary variables  $\Omega$  corresponding to a given discrete index. Crucially, this also works the other way around, yielding the discrete index that corresponds to the frequency closest to a given continuous frequency. This is needed for interpolations, discussed in Sec. II G 7.

An instance of the `FrequencyGrid` class is instantiated in every instance of one of the correlation function classes to parameterize their respective frequency dependencies. The vertex classes naturally require up to three instances of the `FrequencyGrid` each.

The frequency grids are rescaled during mFRG flow calculations, which use the hybridization flow scheme (see Sec. III C). The `FrequencyGrid` class provides all the functionality required for that purpose.

As a side note, two alternative frequency grids have been implemented. One is a hybrid grid, which consists of a quadratic part at small frequencies, a linear part at intermediate frequencies, and a rational part at large frequencies. The other uses polar coordinates to parametrize the two-dimensional frequency dependence of three-point functions, i.e., the  $K_2$  and  $K_2'$  classes. Which grid is to be used is controlled by the `GRID` flag (see Sec. II I). In our experience, the non-linear grid explained at the beginning of this section is the most useful if the scale parameters  $A$  are chosen suitably.

## 5. Frequency integration

The following passage is taken almost verbatim from the Ph.D. thesis of E. Walter.<sup>28</sup>

Computing numerical integrals with high accuracy is a crucial ingredient for obtaining correct results in the context of the diagrammatic calculations discussed here. At the same time, the integrator is also critical for the performance of the computation, since evaluating integrals constitutes the computationally most expensive part of the code. For these reasons, we use an adaptive integration routine that automatically determines where to evaluate the integrand within the integration domain. Regions with sharp features require many evaluation points in order to get high accuracy,

while in regions where the integrand is smooth, fewer evaluations suffice, which increases the performance of the computation. Such an adaptive integrator is really indispensable for the problem at hand. Non-adaptive routines like a simple trapezoidal or Simpson rule on an equidistant grid often lead to systematically wrong results.

We use  $n$ -point integration rules that approximate integrals of the kind  $\int_a^b F(x) dx \approx \sum_{j=1}^n F(x_j) w_j$  with nodes  $x_j$  and corresponding weights  $w_j$ . The integrator we use and which is implemented in the `Adapt` class in the code is an adaptive 4-point Gauss–Lobatto routine with a 7-point Kronrod extension and a 13-point Kronrod extension as an error estimate, as detailed in Ref. 53. The benefit of Gauss–Lobatto rules, compared to, e.g., the widely used Gauss–Kronrod rules, is that the nodes include the endpoints of the integration domain. This allows us to subdivide the domain at the nodes of the integration rule and reuse points that have been computed previously, which is preferential in terms of performance. Similarly, the Kronrod extensions of a Gauss–Lobatto rule reuse all points from a corresponding lower-point rule and simply add additional points, which effectively allows us to get two different rules from one set of evaluation points.

The nodes  $x_j$  of the 4-point Gauss–Lobatto rule with 7-point and 13-point Kronrod extensions are distributed as shown in Fig. 2. There, the lower row indicates the values of the nodes for integration boundaries  $a = -1$ ,  $b = 1$  (for other values of  $a$ ,  $b$ , the values have to be rescaled correspondingly). The four-point Gauss–Lobatto rule (GL4) and four-point Gauss–Lobatto with seven-point Kronrod extension (GLK7) use the following points:

$$\text{GL4}(x_0, x_6) = \sum_{j \in \{0,2,4,6\}} F(x_j) w_j, \quad (14a)$$

$$\text{GLK7}(x_0, x_6) = \sum_{j=0}^6 F(x_j) w_j. \quad (14b)$$

The smaller marks between the nodes  $x_0, \dots, x_6$  in the graphical representation above indicate the additional 6 points that are added in the 13-point Kronrod extension (GLK13), which are only known numerically (these and the weights  $w_j$  are found in Ref. 53).

The recursive algorithm of the integrator then works as shown in Fig. 3. Note that the error estimate  $I_s$  is determined only once for the full integral and then reused for each subinterval in order to avoid infinite recursions in subintervals. A typical recommended value for the relative accuracy is  $\epsilon = 10^{-5}$ , which is set by the global variable `integrator_tol` (see Table II).

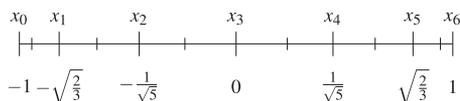


FIG. 2. Distribution of the nodes  $x_j$  of the 4-point Gauss–Lobatto rule with 7-point and 13-point Kronrod extensions. The lower row indicates the values of the nodes for integration boundaries  $a = -1$ ,  $b = 1$ .

Error estimate:  $I_s = \text{GLK13}(a, b)$

$$I_1 = \text{GL4}(a, b) \\ I_2 = \text{GLK7}(a, b)$$

$$|I_2 - I_1| < \epsilon \cdot I_s \xrightarrow{\text{yes}} \int_a^b F(x) dx = I_2$$

$$\text{no} \rightarrow x_0 \leftarrow a, x_6 \leftarrow b$$

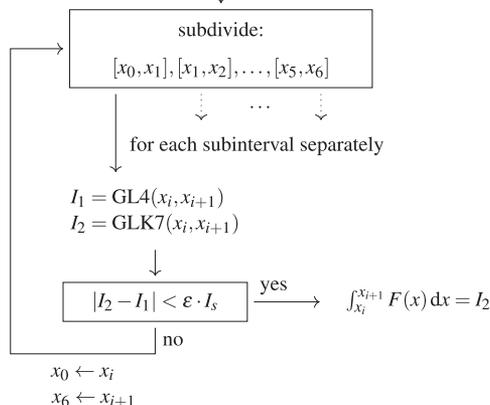


FIG. 3. Schematic illustration of the integration algorithm: an adaptive 4-point Gauss–Lobatto routine with a 7-point Kronrod extension and a 13-point Kronrod extension as an error estimate.

## 6. Asymptotic corrections to frequency integrals

In Sec. II G 5, it was explained how frequency integrations over a finite interval  $[a, b]$  are performed. Since diagrammatic calculations require integrations over the full frequency axis (or summations over an infinite set of discrete Matsubara frequencies for the finite- $T$  MaF), the contributions to the integral resulting from the high-frequency asymptotics of the integrands have to be treated as well. This is particularly relevant for slowly decaying integrands, which occur often, as the correlation functions arising in the present context typically only decay as  $\sim 1/\nu$  or  $\sim 1/\nu^2$ .

In the KF and the zero- $T$  MaF, involving continuous frequency integrations, a naive treatment turned out to be sufficient: Since the frequency axes are discretized non-uniformly, as described in Sec. II G 4, the largest discrete frequency grid point is always so large that the high-frequency tails can be treated via quadrature, ignoring the minuscule contributions of even larger frequencies. For finite- $T$  MaF computations, which involve infinite sums, the code provides two options for the treatment of high-frequency tails in the integrand: (i) The tails can be treated via quadrature by approximating the sum with an integral and then following the same logic as in the KF. (ii) For bubble computations, the lowest order contribution from the bare bubble, which is known analytically, can be used. This is justified by the fact that in the high-frequency asymptotic limit, the non-trivial contributions due to interactions encoded in the self-energy have decayed, and only the bare contribution is responsible for the asymptotic behavior. The first or second

option is chosen with the `ANALYTIC_TAILS` parameter flag (see also Sec. II I).

## 7. Interpolation routines

Whenever the value of a correlation function at some continuous frequency argument is required, in particular during frequency integrations, the data stored on discrete frequency grids has to be interpolated. In addition, the diagrammatic algorithms discussed here have feedback between the three two-particle channels, which all have their own channel-dependent parametrizations. This necessitates accurate interpolations between different frequency parametrizations; otherwise, errors accumulate over the course of a computation.

To handle the interpolation of multidimensional correlation functions, we implemented multilinear interpolation and cubic spline interpolation using cubic Hermite splines. While spline interpolation is robust against minor inaccuracies of the data points and offers faster convergence in the number of frequency points for smooth functions, multilinear interpolation is generally faster numerically. Having tried out both options, we prefer linear interpolation, as spline interpolation only really becomes useful for better precision if the function is already well resolved.

Regarding linear interpolation, the code offers options: One can either interpolate on the grid of frequencies  $\nu$  or on the grid of auxiliary frequencies  $\Omega$ , which are equidistantly spaced on the interval  $[-1, 1]$  (see Sec. II G 4). We found the latter option to be more accurate. The global parameter `INTERPOLATION` specifies which type of interpolation shall be used (see also Sec. II I).

## 8. Data structures

The central low-level data structure used for storing and retrieving numerical data inside the code is the `dataBuffer` class. It was devised with the two main intentions of efficiency and flexibility in mind (see also our discussion of the main design choices for the codebase in the introduction, Sec. I). On the one hand, it should enable building integrands that return scalar- or vector-valued entries as efficiently as possible, particularly avoiding conditional (“if-else”) statements during runtime, as these prevent optimizations such as loop-vectorization or function inlining. On the other hand, it should be useable in all parts of the codebase, e.g., for both calculations with interpolations on continuous frequency grids and for finite- $T$  MaF calculations, which only require indexing of discrete data points.

The `dataBuffer` class is structured as follows. It builds upon the `dataContainerBase` class, which is used to represent multi-dimensional tensors, allowing scalar and vector-valued access to contiguous elements. The `DataContainer` class then inherits `dataContainerBase`, adding frequency information. It contains a multi-dimensional frequency grid (see Sec. II G 4) to parameterize all its associated frequency arguments and provides functions to analyze the resolution of frequency grids. Inheriting the `DataContainer` class, the `Interpolator` classes then implement the different interpolation routines outlined in Sec. II G 7. Multilinear cubic spline interpolations require pre-computation and storage of interpolation coefficients, whereas linear interpolations happen on the fly. Finally, the `dataBuffer` class inherits both the

`Interpolator` and the `DataContainer` classes and can be used in actual computation. In addition, it can update and optimize grid parameters as required.

## 9. Template arguments

Another performance-critical aspect of the codebase is its heavy use of templates. In particular, the propagation of template arguments as specified by preprocessor flags enables the determination of the required diagrammatic combinations for any given computation at compile time. Selecting and combining the necessary vertex contributions this way, e.g., for contributions to specific asymptotic classes, enables further optimization by the compiler. However, the ubiquity of template arguments comes at the expense of readability in many places.

## H. Tests

The code includes a large number (178 as of writing) of self-explanatory unit tests that run checks on the low-level parts of the codebase. They are implemented using the popular `Catch2` library<sup>34</sup> and are invoked from a separate C++ source file, `unit_tests.cpp`, which should be built separately from the main source file. From inside this file, more involved and expensive tests can be started if desired. These include detailed tests of the ODE solver or perturbation theory, which are too expensive to be part of the unit test suite. Finally, the code includes functionality to produce reference data that can be used later to compare the results of a calculation after changes to the code have been made. We have found it immensely useful to include many unit tests in the codebase, as they can tell almost immediately if a single technical part of the code has broken. Moreover, having a way to compare the results of very involved computations that involve large parts of the codebase at once is useful to catch logical errors. We wholeheartedly recommend both.

## I. Parameters

Before any individual calculation can be started, a number of parameters have to be set. As the code provides a large degree of flexibility, the number of possible parameter choices is large. Most of these parameters are set inside the corresponding header files before compilation. The reason for this is that, depending on these choices, often different functionality of the code is invoked, depending, e.g., on the choice of formalism. This is achieved by defining preprocessor macros accordingly, which makes the corresponding functionality accessible. As discussed previously in too many preprocessor macros (“flags”), while this approach was useful for implementing new functionality quickly, in the long run, it turned out to be problematic with regard to the readability and maintainability of the code. Table I provides a list (albeit incomplete) of the most important preprocessor flags used in the code with a short description of each.

In addition, global parameters have to be set, which specify settings like the resolution of the frequency grid, convergence criteria, or start- and end-points of an mFRG flow. Table II provides a non-exhaustive list of those.

Finally, it should be mentioned that once the code has been compiled and the resulting executable is to be called, it requires three run-time arguments: The first one invokes an mFRG run if it

TABLE I. Incomplete list of the most important preprocessor macros to be set before compilation.

Macro name	Possible values	Description
ADAPTIVE_GRID	...	If defined, use the optimization routine to find the best scale factor $A$ of the frequency grid; if undefined, just rescale the grid. Warning: Can be expensive and unreliable in the presence of numerical artifacts
ANALYTIC_TAILS	0, 1	0 for false; 1 for true. If true, the analytic expression for the bare bubble is used to treat the high-frequency asymptotics during bubble computations in the finite- $T$ MaF
BARE_SE_FEEDBACK	...	If defined, only bare selfenergy is used. It only makes sense if <code>STATIC_FEEDBACK</code> is defined. Useful for benchmarks with previous Keldysh fRG schemes
CONTOUR_BASIS	0, 1	0 for false, 1 for true: If true, no Keldysh rotation is performed, and the contour basis is used instead to parameterize the Keldysh components of all correlation functions. It is useful for comparisons with results that use this convention. Not as well tested and, therefore, not recommended for production runs
DEBUG_SYMMETRIES	0, 1	0 for false; 1 for true. Performs computations without the use of symmetries, if true. Useful for debugging purposes
GRID	0, 1, 2	Controls which frequency grid is to be used. 0 for the non-linear grid, 1 for the hybrid grid, and 2 for the polar grid. Recommendation: 0. See also Sec. II G 4
KATANIN KELDYSH_FORMALISM	...	If defined, the Katanin extension is used during fRG computations Determines whether calculations shall be performed in the Keldysh or Matsubara formalism. 0 for Matsubara formalism (MaF); 1 for Keldysh formalism (KF)
MAX_DIAG_CLASS	1, 2, 3	Defines the diagrammatic classes that will be considered: 1 for only $K_1$ , 2 for $K_1$ and $K_2$ , and 3 for the full dependencies. Useful for debugging purposes and for computations in second-order perturbation theory, or if <code>STATIC_FEEDBACK</code> is defined, when only $K_1$ is required
NDEBUG	...	If defined, assert functions are switched off. Recommended setting for production runs
PARTICLE_HOLE_SYMM	0, 1	0 for false; 1 for true. If true, particle-hole symmetry is assumed
PT2_FLOW	...	If defined, only compute the flow equations up to $O(U^2)$ . Only makes sense for pure $K_1$ calculations. It is useful as a consistency check together with independent PT2 calculations
REG	2, 3, 4, 5	Specifies the mFRG flow regulator to be used. 2: $\Delta$ -flow, 3: $\omega$ -flow, 4: $U$ -flow, 5: $T$ -flow. For details, see Sec. III C 2
REPARAMETRIZE_FLOWGRID	...	If defined, the flow parameter is reparametrized according to Sec. III C 1. Only recommended for the $\Delta$ -flow
SBE_DECOMPOSITION	0, 1	0 for false; 1 for true. If true, the SBE decomposition is used to parameterize the vertex and the flow equations. Only implemented in the MaF!
SELF_ENERGY_FLOW_CORRECTIONS	0, 1	0 for false; 1 for true. If true, corrections to the flow equations for the vertex from the self-energy, starting at $\ell = 3$ , are included
STATIC_FEEDBACK	...	If defined, use static $K_1$ inter-channel feedback as performed in 11. Only makes sense for pure $K_1$ calculations
SWITCH_SUM_N_INTEGRAL	0, 1	0 for false; 1 for true. If true, the sum over internal Keldysh indices is performed before the frequency integration. Recommended setting: 1
USE_ANDERSON_ACCELERATION	0, 1	0 for false; 1 for true. If true, Anderson acceleration is used to converge parquet iterations and self-energy iterations in mFRG faster
USE_MPI	...	If defined, MPI is used for parallelization across multiple nodes
USE_SBEb_MFRG_EQS	0, 1	Determines which version of the SBE approximation shall be used. 0 for SBEa, 1 for SBEb. Only implemented in the MaF!
VECTORIZED_INTEGRATION	0, 1	0 for false; 1 for true. If true, integrals are performed with vector-valued integrands. For Keldysh, vectorization over Keldysh indices. For Matsubara at finite $T$ , vectorization over the Matsubara sum
ZERO_TEMP	0, 1	0 for false; 1 for true. If true, temperature $T = 0$ is assumed

TABLE II. Incomplete list of global parameters to be set before compilation.

Parameter name	Type	Description
converged_tol	double	Tolerance for loop convergence in mFRG
COUNT	int	Used to set the number of frequency points in the MaF. For details, see the definitions in the file <code>frequency_parameters.hpp</code>
Delta_factor_K1	int	Scale factor for the frequency grid of the K1 vertex class
Delta_factor_SE	int	Scale factor for the frequency grid of the self-energy
Delta_factor_K2_w	int	Scale factor for the frequency grid of the bosonic frequency of the K2 and K2' vertex classes
Delta_factor_K2_v	int	Scale factor for the frequency grid of the fermionic frequency of the K2 and K2' vertex classes
Delta_factor_K3_w	int	Scale factor for the frequency grid of the bosonic frequency of the K3 vertex class
Delta_factor_K3_v	int	Scale factor for the frequency grid of the fermionic frequencies of the K3 vertex class
EQUILIBRIUM	bool	If true, use equilibrium FDRs for propagators
glb_mu	double	Chemical potential – w.l.o.g. ALWAYS set to 0.0 for the AM!
integrator_tol	double	Integrator tolerance
inter_tol	double	Tolerance for closeness to grid points when interpolating
INTERPOLATION	linear, linear_on_aux, cubic	Interpolation method to be used. <code>linear</code> : linear interpolation on the frequency grid. <code>linear_on_aux</code> : linear interpolation on the grid for the auxiliary frequency $\Omega$ . <code>cubic</code> : Interpolation with cubic splines (warning: expensive!)
Lambda_ini	double	Initial value of the regulator $\Lambda$ for an mFRG flow
Lambda_fin	double	Final value of the regulator $\Lambda$ for an mFRG flow
Lambda_scale	double	Scale of the log substitution, relevant in the hybridization flow
dLambda_initial	double	Initial step size for ODE solvers with adaptive step size control
nBOS	int	Number of bosonic frequency points for the $K_1$ vertex class
nFER	int	Number of fermionic frequency points for the self-energy
nBOS2	int	Number of bosonic frequency points for the $K_2$ and $K_2'$ vertex classes
nFER2	int	Number of fermionic frequency points for the $K_2$ and $K_2'$ vertex classes
nBOS3	int	Number of bosonic frequency points for the $K_3$ vertex class
nFER3	int	Number of fermionic frequency points for the $K_3$ vertex class
U_NRG	std::vector<double>	Vector with the values of $U$ in units of $\Delta$ that an mFRG flow should cover. Serve as checkpoints for the flow. It is useful for benchmarking purposes if data from other methods at precise parameter points are available
VERBOSE	bool	If true, detailed information about all computational steps is written into the log file. Recommended setting for production runs: <code>false</code>
nmax_Selfenergy_iterations	int	Maximal number of self-energy iterations to be performed during an mFRG flow for $\ell \geq 3$ . Default value: 10
tol_selfenergy_correction_abs	double	Absolute tolerance for self-energy iterations in mFRG. Default value: $10^{-9}$
tol_selfenergy_correction_rel	double	Relative tolerance for self-energy iterations in mFRG. Default value: $10^{-5}$

is a positive integer, specifying the maximal number of loop orders calculated during the mFRG flow. Alternatively, if it is set to 0 or  $-1$ , a parquet or PT2 calculation is started, respectively. The second is a positive integer and specifies the number of nodes to be utilized. The third runtime argument defines the temperature for the

calculation and was introduced to easily enable parameter sweeps without having to recompile the code every time. Note that its value is irrelevant for calculations that have the flag `ZERO_TEMP` set to 1 or if an mFRG run is performed with the flag `REG` set to 5, which employs the temperature flow.

### III. ALGORITHMS

In the third main part of the paper, we finally describe three diagrammatic algorithms that have been implemented. These are second-order perturbation theory (PT), a self-consistent solution of the parquet equations, and the flow equations provided by the multiloop functional renormalization group (mfRG). For all three methods, we first give some theoretical background before describing schematically how the algorithms are implemented and what functions are being used.

#### A. Perturbation theory

The simplest computations that can be performed with the code are perturbation theory calculations. While these are easy to implement in the second order, going to higher orders involves an increasing number of diagrams, which can in principle be evaluated separately. This is, however, not always straightforward, e.g., if symmetries are to be exploited: individual diagrams of the perturbation series do not all have the same symmetries as a full vertex, such that symmetry-related diagrams have to be provided, which can become tedious. Alternatively, the flag `DEBUG_SYMMETRIES` can be set to 1, see Sec. II 1, in which case the code does not attempt to exploit symmetries. As higher-order perturbation theory has so far only been performed for testing purposes and consistency checks (see, e.g., Chap. 7 in Ref. 28), we refrain from going into further detail here. Instead, we focus just on the second-order case and on Hartree-Fock theory for the self-energy relevant to the aAM.

##### 1. Hartree-Fock

As elaborated in Ref. 1, it is helpful to replace the bare propagator  $G_0$  by the Hartree-propagator  $G_H$ , which is shifted by the Hartree-term of the self-energy,

$$G_0^R \rightarrow G_H^R = \frac{1}{v - \varepsilon_d + i\Delta - \Sigma_H^R}. \quad (15)$$

For the sAM, this is almost trivial, as the retarded component of the Hartree term reads  $\Sigma_H^R = U/2$ , which simply yields  $G_H^R = (v + i\Delta)^{-1}$ . For the aAM, on the other hand, the Hartree-term can be computed self-consistently.

For this purpose, the class `Hartree_Solver` provides the function `compute_Hartree_term_bracketing`. It computes  $\Sigma_H^R$  via

$$\Sigma_H^R = U \int \frac{d\nu'}{2\pi i} G_H^<(\nu'), \quad (16)$$

where in thermal equilibrium, the relation  $G^<(\nu) = -2i n_F(\nu) \text{Im} G^R(\nu)$  is used with the Fermi function  $n_F(\nu) = 1/(1 + e^{\nu T})$ . As  $\Sigma_H^R$  enters both sides of Eq. (16), this calculation is performed self-consistently using a simple bracketing algorithm.

In addition, the `Hartree_Solver` class provides the function `compute_Hartree_term_oneshot`, which evaluates Eq. (16) just once, given a provided self-energy for  $G^R(\nu)$ . This function is invoked in the context of parquet iterations and evaluations of mfRG flow equations to update the Hartree term of the aAM.

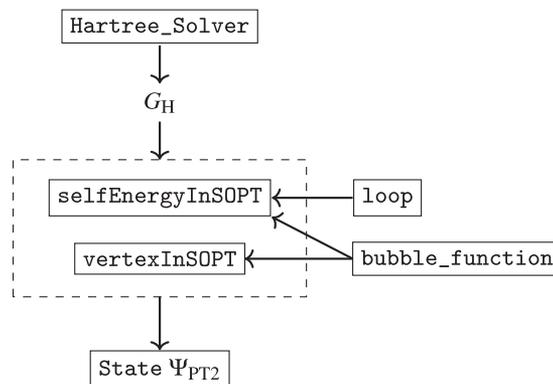


FIG. 4. Schematic depiction of the function `sopt_state`.

Finally, the `Hartree_solver` class provides functionality to check the fulfillment of the Friedel sum rule<sup>55</sup>  $\langle n_\sigma \rangle = \frac{1}{2} - \frac{1}{\pi} \arctan [(\varepsilon_d + \Sigma(0))/\Delta]$ , which the self-consistent Hartree term fulfills at  $T = 0$ .

#### 2. Second order perturbation theory (PT2)

The self-energy and vertex in second-order perturbation theory are computed via the function `sopt_state`, which works as depicted in Fig. 4. It first initializes a bare state (see Sec. II C 3), given the system parameters and the current value of the regulator  $\Lambda$ . For the aAM, this already includes a self-consistent calculation of the Hartree term (see Sec. III A 1). Then, it invokes the function `selfEnergyInSOPT`, which computes the single diagram for the dynamical part of the self-energy in PT2 by first computing a bare bubble in the  $a$ -channel using the `bubble_function` (see Sec. II D 1), with two bare vertices, and then closing the loop over that bare bubble with the Hartree-propagator using the `loop` function (see Sec. II D 2).

Thereafter, the vertex is computed using the function `vertexInSOPT`, which simply invokes the `bubble_function` three times, once for each of the three two-particle channels  $a$ ,  $p$ , and  $t$ , using two bare vertices, adding each result to the vertex.

In total, this procedure yields all diagrams for the dynamical part of the self-energy and the vertex in PT2, using the Hartree-propagator  $G_H$ . For the precise diagrammatic definition of PT2, see Appendix F in Ref. 1.

#### B. Parquet equations

The parquet formalism<sup>56</sup> provides a self-consistent set of equations for the self-energy  $\Sigma$  and the three two-particle reducible vertices  $\gamma_r$  with  $r \in \{a, p, t\}$ . The latter are given by the Bethe-Salpeter equations (BSEs)

$$\gamma_a = I_a \Gamma, \quad (17a)$$

$$\gamma_p = \frac{1}{2} I_p \Gamma, \quad (17b)$$

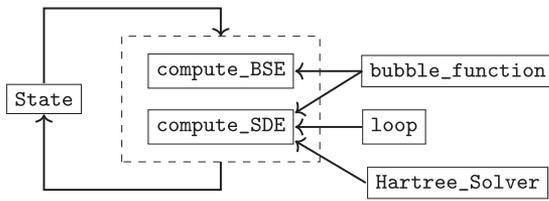


FIG. 5. Schematic depiction of the `parquet_solver` function.

$$\gamma = - \left( \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \right) = - \left( \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \right) \quad (17c)$$

where  $I_r = \Gamma - \gamma_r$  is the two-particle irreducible vertex in channel  $r$ . The self-energy is given by the Schwinger–Dyson equation (SDE),

$$\Sigma = - \left( \text{---} \right) - \frac{1}{2} \left( \text{---} \right) \quad (18)$$

which includes the Hartree term discussed in Sec. III A 1. Together, these equations close once the fully irreducible vertex  $R$  is provided, for example, by employing the PA, as discussed in Sec. I B.

In practice, these equations are solved iteratively. The code provides functions to evaluate the right-hand sides of the BSEs and the SDE, called `compute_BSE` and `compute_SDE`. Schematically, the parquet solver works as depicted in Fig. 5. Inside the code, a parquet computation is started by the function `run_parquet`. It first initializes a state using PT2, as detailed in Sec. III A 2, before the `parquet_solver` function is called. Internally, the `parquet_solver` calls `parquet_iteration`, which evaluates the BSEs and the SDE, given a provided input state, and combines them into an output state. The corresponding functions `compute_BSE` and `compute_SDE` use the machinery described in Secs. II C and II D to evaluate Eqs. (17) and (18). In practice, symmetrizing Eq. (17), i.e., computing the sum of the right-hand side as is and with  $I_r$  and  $\Gamma$  interchanged and dividing by two, has proven beneficial for stability.

In addition, we found it helpful to combine all three ways to evaluate the SDE, Eq. (18) (see Appendix D in Ref. 1).

The `parquet_solver` can either proceed directly from one iteration to the next, or it can combine multiple results from previous iterations using mixing schemes to improve convergence. For example, one can combine the two most recent iterations with a mixing factor as outlined in Eq. (G4) of Ref. 1. One may start with a mixing factor of around 0.5, which can be reduced automatically if the convergence properties of the calculation are poor. In addition, one can use Anderson acceleration<sup>57,58</sup> to combine multiple previous iterations for a prediction of the next iteration. We have found that this leads to faster convergence in the vicinity of the solution but does not extend the parameter range where convergence can be reached.

The parquet solver can also be used for calculations in the random phase approximation (RPA). Switching off the BSEs in two of the three two-particle channels readily yields the RPA-ladder in the other channel.

### C. mFRG

In fRG,<sup>59</sup> the self-energy and vertex are interpolated between the initial and final values of a single-particle parameter  $\Lambda$  introduced into the bare propagator  $G_0$ . The initial value  $\Lambda = \Lambda_i$  should be chosen such that the theory is solvable at that point; in practice, it typically suffices that very good approximations of  $\Sigma^{\Lambda_i}$  and  $\Gamma^{\Lambda_i}$  can be obtained by PT2 or by converging the parquet equations. The fRG then provides a set of differential “flow” equations in  $\Lambda$  for  $\Sigma^\Lambda$  and  $\Gamma^\Lambda$ , which yield the final results  $\Sigma^{\Lambda_f}$  and  $\Gamma^{\Lambda_f}$  at the actual point of interest  $\Lambda = \Lambda_f$ . In the multiloop fRG framework, these flow equations are derived from the parquet equations by differentiation with respect to the flow parameter  $\Lambda$ , as detailed in Ref. 18. This yields an infinite set of contributions of increasing “loop order”  $\ell$ ,

$$\dot{\Gamma} = \sum_{r \in \{a,p,t\}} \dot{\gamma}_r \quad (19a)$$

$$\dot{\gamma}_r = \sum_{\ell=1}^{\infty} \dot{\gamma}_r^{(\ell)}, \quad (19b)$$

where a dot represents a derivative with respect to  $\Lambda$ . Diagrammatically, the  $\ell$ -loop contributions in the  $a$  channel read

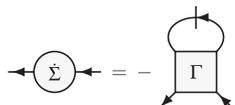
$$\dot{\gamma}_a^{(1)} = \left( \text{---} \right) \quad (20a)$$

$$\dot{\gamma}_a^{(2)} = \underbrace{\left( \text{---} \right)}_{\dot{\gamma}_{a,L}^{(2)}} + \underbrace{\left( \text{---} \right)}_{\dot{\gamma}_{a,R}^{(2)}} \quad (20b)$$

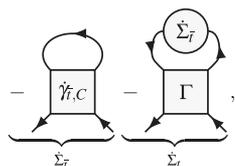
$$\dot{\gamma}_a^{(\ell+2)} = \underbrace{\left( \text{---} \right)}_{\dot{\gamma}_{a,L}^{(\ell+2)}} + \underbrace{\left( \text{---} \right)}_{\dot{\gamma}_{a,C}^{(\ell+2)}} + \underbrace{\left( \text{---} \right)}_{\dot{\gamma}_{a,R}^{(\ell+2)}} \quad (20c)$$

and analogously in the other two channels  $p$  and  $t$ . Here,  $\gamma_r^{(\ell)} = \sum_{r' \neq r} \gamma_{r'}^{(\ell)}$ , and Eq. (20c) applies for all higher loop orders  $\ell + 2 \geq 3$ . The double-dashed bubble in Eq. (20) corresponds to a sum of two terms,  $\dot{\Gamma} = \dot{G}G + G\dot{G}$ , where  $\dot{G} = S + G\dot{\Sigma}G$  with the single-scale propagator  $S = \partial_\Lambda G|_{\Sigma=\text{const}}$  and the Katanin substitution.<sup>48</sup>

The multiloop flow equation for the self-energy reads



$$\Sigma = - \text{bubble}(\Gamma) \quad (21a)$$



$$\dot{\gamma}_{i,C} = - \text{bubble}(\dot{\Sigma}_r) - \text{bubble}(\dot{\Sigma}_t) \quad (21b)$$

with  $\dot{\gamma}_{i,C} = \sum_\ell (\dot{\gamma}_{a,C}^{(\ell)} + \dot{\gamma}_{p,C}^{(\ell)})$ , where the single-dashed line denotes the single-scale propagator  $S$  from above.

Historically, fRG flow equations have been derived from a generating functional, yielding an exact hierarchy of flow equations which couple  $n$ -point vertices of increasing order.<sup>60</sup> As the six-point vertex, which contributes to the flow equation of  $\Gamma$ , see Eq. (19) in Ref. 1, is inaccessible numerically, its contribution is often neglected completely, resulting in the so-called “one-loop” flow equations. This, however, results in an unphysical dependence of the final result of the flow on the choice of regulator (as the flow equations no longer constitute total derivatives) and also introduces a bias toward ladder diagrams.<sup>16,61</sup>

The multiloop framework builds upon the one-loop scheme by iteratively adding precisely those two-particle reducible diagrammatic contributions to the flow equations that are required to reinstate total derivatives with respect to  $\Lambda$  and thereby reproduce the solution of the parquet equations. In that sense, it provides an alternative scheme for solving the parquet iterations via differential equations. From a computational standpoint, the mfRG flow equations introduce a complication compared to the one-loop flow equations, in that the right-hand sides of the flow equations for both  $\Gamma$  and  $\Sigma$  involve the differentiated self-energy and vertex. In order to still be able to use standard algorithms for ordinary differential equations, a scheme was outlined in Ref. 16 to include those differentiated quantities iteratively. Starting from the one-loop term Eq. (21a) to evaluate the flow Eq. (20) for  $\Gamma$ , these are then iterated with the multiloop corrections (21b) at every step of the flow until convergence is reached. The number of iterations required for convergence at this point can again be reduced using Anderson acceleration, as described in Sec. III B.

From the code, an mfRG-flow computation can be started with the function `n_loop_flow`, which requires only the string for the name of the output file and a set of parameters. It is overloaded to enable checkpointing, i.e., it is possible to continue a previously started computation from a given iteration. This is particularly useful for demanding jobs that take a long time, and it is highly recommended to any user.

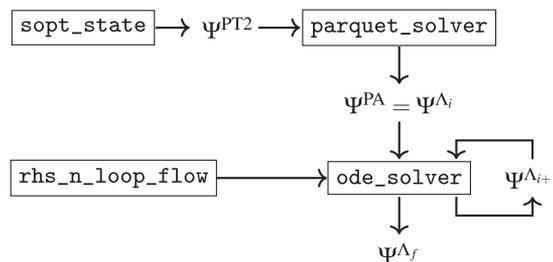
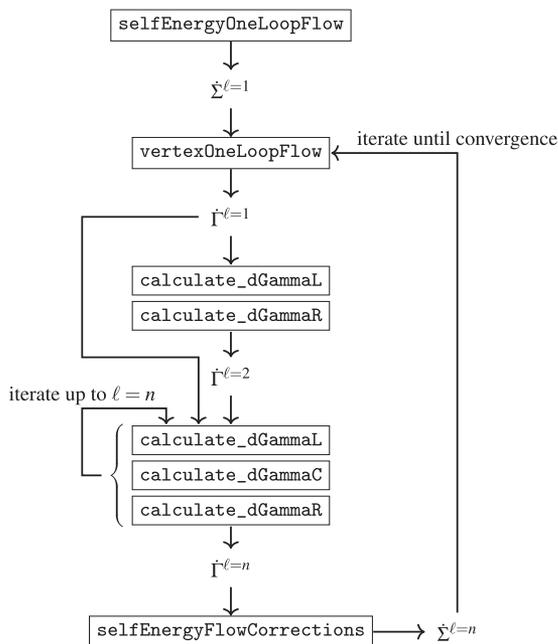


FIG. 6. Schematic depiction of the function `n_loop_flow`.

The function `n_loop_flow` works as shown in Fig. 6. It first initializes a state using PT2 with the function `sopt_state`, see Sec. III A 2, and then uses this result as a seed for a full parquet computation at the initial value of the regulator  $\Lambda_i$  with the `parquet_solver` function, see Sec. III B. This provides a suitable starting point for the following mfRG calculation.

The `ode_solver` function carries out the actual calculation of solving the mfRG flow. It uses an instance of the `rhs_n_loop_flow_t` class, which provides a wrapper to the function `rhs_n_loop_flow`, which in turn evaluates the right-hand side of the flow equations given an input state at a given value of  $\Lambda$ . This is performed iteratively by loop order according to flow Eq. (20), including self-consistent iterations for the self-energy starting at the three-loop level, as outlined earlier. The function `rhs_n_loop_flow` is structured as shown in Fig. 7. Starting from the self-energy and vertex from the previous step of the ODE-solver, it evaluates the right-hand sides of the flow equations by first computing the one-loop term of the flow equation for the self-energy, Eq. (21a), with the function `selfEnergyOneLoopFlow`. The result is then used to evaluate the one-loop term of the flow equation for the two-particle reducible vertices  $\dot{\gamma}_r$  and Eq. (20), involving a fully differentiated bubble. Then, the one-loop result is used to evaluate the two-loop contribution, Eq. (20b), which consists of two terms: one where the differentiated one-loop contribution  $\dot{\gamma}_r^{(1)}$  is used as the left part of a bubble contraction with the full vertex, and one where it is used on the right side. These two terms are computed using the functions `calculate_dGammaL` and `calculate_dGammaR`, respectively. Next, the three-loop contribution is computed, which involves both the one-loop and the two-loop results [see Eq. (20c)]. Again, the functions `calculate_dGammaL` and `calculate_dGammaR` are invoked, and in addition, the function `calculate_dGammaC` is invoked to compute the “center term” involving two bubble contractions of  $\dot{\gamma}_r^{(1)}$  with the full vertex, once to the left and once to the right. As the structure of the flow equations does not change from this point on, this part is iterated until the maximally desired loop number  $n$  (which is given as a runtime parameter; see Sec. II 1) is reached. The resulting center terms of the  $a$  and  $p$  channels are then used to evaluate the multiloop corrections to the self-energy, according to Eq. (21b). This updates the differentiated bubble used in the computation of the one-loop terms  $\dot{\gamma}_r^{(1)}$ , such that the whole process is finally iterated from that point on until convergence is reached, as determined by the parameters `tol_selfenergy_correction_abs` and



**FIG. 7.** Structure of the function `rhs_n_loop_flow`, including multiloop iterations up to loop order  $\ell = n$  and self-consistent self-energy iterations due to the multiloop corrections.

`tol_selfenergy_correction_rel` (see Sec. II I). All functions invoked by `rhs_n_loop_flow`, of course, make heavy use of the main functionality outlined in Sec. II D.

As a side note, it is possible to parameterize the vertex using the single-boson exchange (SBE) decomposition<sup>62–67</sup> and to rewrite the mFRG flow equations in this language, as outlined in Ref. 68. This is achieved by setting the flag `SBE_DECOMPOSITION` to 1. Two versions of the SBE approximation can be used, known as “SBEa” and “SBEb” in the literature.<sup>69</sup> Which version is to be used is controlled by the flag `USE_SBEb_MFRG_EQS` (see Sec. II I). This functionality is, so far, only implemented in the MaF. We, therefore, refrain from providing further details here.

In the final two parts of this section, we discuss the ODE-solver and the different flow schemes.

### 1. Details on the ODE-solver

To solve the mFRG flow equations accurately, a Cash–Karp routine<sup>70</sup> is implemented, which constitutes a fourth-order Runge–Kutta solver with adaptive step size control. An adaptive step-size control is crucial for obtaining accurate results and is hereby strongly recommended for solving fRG flow equations precisely. For a good first guess of the step size in the  $\Delta$ -flow (see Sec. III C 2 a), the flow parameter is reparametrized as  $\Lambda(t) = 5t|t|/\sqrt{1-t^2}$ . For equidistant  $t$ , this parametrization provides large steps for large  $\Lambda$  and small steps for small  $\Lambda$ . This is sensible in the context of the  $\Delta$ -flow, where  $\Lambda$  is gradually reduced to enter ever more challenging parameter regimes.

### 2. Flow schemes

In fRG, one chooses a regulator introduced into the bare propagator  $G_0 \rightarrow G_0^\Lambda$ , i.e., the flow scheme. While the solution of a truncated set of fRG flow equations will depend on this choice, a converged multiloop flow will not, as it reproduces the self-consistent solution of the parquet equations. It is generally advisable to choose the most convenient flow scheme for the problem at hand. In particular, the fRG flow can be used to compute a full parameter sweep in one go by choosing a physical parameter as the regulator. Compared to direct solutions of the parquet equations, which have to be computed individually at every point in parameter space, this makes mFRG computations more economical, provided they can be quickly converged in the loop order. In the following, we outline the flow schemes that have been implemented and can be used by setting the `REG` flag and the `Lambda_ini` and `Lambda_fin` parameters accordingly (see Tables I and II).

*a.  $\Delta$ -flow.* The hybridization flow<sup>11</sup> uses  $\Delta$  as the flow parameter, starting at a very large value and decreasing  $\Delta$  to a smaller value, keeping the other parameters  $U$  and  $T$  fixed. The hybridization flow thus performs a parameter sweep in  $U/\Delta$  for fixed  $T/U$ . The Keldysh fRG single-scale propagator reads

$$S^R(\nu) = \partial_\Delta G^R(\nu)|_{\Sigma=\text{const}} = -i[G^R(\nu)]^2.$$

In practice, we start the fRG flow from a solution of the parquet equations at large  $\Delta$  (small  $U/\Delta$ ), where that solution can be easily obtained. For historical reasons, the hybridization flow is implemented as

$$G_\Lambda^R(\nu) = \frac{1}{\nu - \varepsilon_d + i(\Gamma + \Lambda)/2 - \Sigma_\Lambda^R(\nu)}, \quad (22)$$

inside the code, where  $\Gamma$  is fixed to some arbitrary value and  $\Lambda$  is used to fix the hybridization  $\Delta = (\Gamma + \Lambda)/2$ . Note that keeping  $T/U$  fixed during the  $\Delta$ -flow is a somewhat unconventional choice, as in most works on the AM, the scale  $T/\Delta$  is kept constant. As explained in Ref. 28, keeping  $T/\Delta$  fixed during the  $\Delta$ -flow would lead to additional sharply peaked terms in the single-scale propagator and has hence not been pursued yet.

*b.  $U$ -flow.* An alternative to the  $\Delta$ -flow is the following flow scheme, first introduced in Ref. 71,

$$G_\Lambda^R(\nu) = \frac{\Lambda}{\nu - \varepsilon_d + i\Delta - \Lambda \Sigma_\Lambda^R(\nu)}, \quad (23)$$

starting at  $\Lambda_i = 0$  (or very small, in practice) and flowing toward  $\Lambda_f = 1$ . The corresponding single-scale propagator then reads

$$S^R(\nu) = \partial_\Lambda G^R(\nu)|_{\Sigma=\text{const}} = \frac{\nu - \varepsilon_d + i\Delta}{[\nu - \varepsilon_d + i\Delta - \Lambda \Sigma_\Lambda^R(\nu)]^2}. \quad (24)$$

This flow scheme is called interaction- or  $U$ -flow because increasing  $\Lambda$  effectively amounts to increasing  $U$ . This can be shown by a simple rescaling argument: A bare diagram for  $\Sigma$  (or  $\Gamma$ ) at order  $n$  has  $n$  factors of  $U$  and  $2n - 1$  (or  $2n - 2$ ) factors of  $G_{0,\Lambda}$ , each contributing one factor of  $\Lambda$ . The same scaling behavior in  $\Lambda$  can be achieved

without a  $\Lambda$ -dependent  $G_0$  by multiplying  $U$  with  $\Lambda^2$  and dividing out an extra  $\Lambda$  (or  $\Lambda^2$ ). It hence holds that

$$\Sigma_\Lambda(U) = \Sigma(\Lambda^2 U)/\Lambda, \quad (25a)$$

$$\Gamma_\Lambda(U) = \Gamma(\Lambda^2 U)/\Lambda^2. \quad (25b)$$

Note that at zero temperature, the two flow schemes discussed so far should be equivalent: For  $T = 0$ , the only energy scales of the AM in the wideband limit are  $U$  and  $\Delta$ , so there is only one external parameter  $U/\Delta$  and it does not matter whether  $U$  is increased or  $\Delta$  is decreased.

Historically, the  $U$ -flow has not been very popular, as it does not regulate IR divergences.<sup>59</sup> Nevertheless, it can be used for the AM. In Ref. 1, we found that, for a truncated 1-loop Keldysh fRG flow at finite  $T$ , this scheme produces inferior results compared to the  $\Delta$ -flow when benchmarked against numerically exact NRG data. Still, the  $U$ -flow has the nice property that it keeps  $T/\Delta$  fixed.

*c. T-flow.* Using temperature as the fRG flow parameter has been popular in the past when performing fRG computations in the MaF.<sup>72,73</sup> It has been argued that temperature cannot be used for this purpose in Keldysh fRG computations,<sup>11</sup> the reason being that a truncated fRG flow does not preserve fluctuation–dissipation relations (FDRs). However, solutions to the parquet equations *do* fulfill the FDRs. If the FDRs are not used explicitly during mFRG calculations (as this would mix FDRs at different temperatures and hence introduce an inconsistency), it should also be possible to obtain these solutions by converging an mFRG flow. Instead of the standard FDR, which relates  $G^K$  and  $G^R$ , in this scheme, the general expression for the Keldysh component of the propagator should be used, which reads<sup>28</sup>

$$G^K(\nu) = G^R(\nu) \left[ \Sigma^K(\nu) - 2i\Delta \tanh\left(\frac{\nu}{2T}\right) \right] G^A(\nu). \quad (26)$$

The Keldysh component of the single-scale propagator is then

$$S^K(\nu) = \partial_T G^K(\nu)|_{\Sigma=\text{const}} = \frac{i\Delta\nu}{T^2 \cosh^2\left(\frac{\nu}{2T}\right)} |G^R(\nu)|^2. \quad (27)$$

Note that its retarded component is zero;  $S^R(\nu) = 0$ , as  $G^R(\nu)$  does not depend explicitly on  $T$ . While preliminary numerical results suggest that this scheme indeed performs well, a systematic study of the temperature flow in Keldysh fRG is left for future work. So far, at the time of writing, the temperature flow described earlier can only be used in the KF; corresponding regulators in the MaF, as in Refs. 72 and 73, have not been implemented.

*d.  $\nu$ -flow.* Using a frequency regulator of the form  $G_{0,\Lambda}(i\nu) = G_0(i\nu)\Theta_\Lambda(i\nu)$  with  $\Theta_\Lambda(i\nu) = \nu^2/(\nu^2 + \Lambda^2)$  has been a popular choice in the literature for (m)fRG calculations in the Matsubara formalism.<sup>74,75</sup> However, in this form, the frequency regulator cannot be used in the Keldysh formalism, as the analytical continuation of  $\Theta_\Lambda(i\nu)$  gives  $\Theta_\Lambda^R(\nu) = \nu^2/(\nu^2 - \Lambda^2 + 2|\nu|i0^+)$  with a branch cut

for  $\nu < 0$ . One can, however, change the form of the regulator to  $\Theta_\Lambda(i\nu) = |\nu|/(|\nu| + \Lambda)$ , for which the retarded counterpart reads

$$\Theta_\Lambda^R(\nu) = \frac{\nu}{\nu + i\Lambda}, \quad (28)$$

which is a well-behaved function. This choice is implemented as

$$G_\Lambda^R(\nu) = \frac{\Theta_\Lambda^R(\nu)}{\nu - \varepsilon_d + i\Delta - \Theta_\Lambda^R(\nu)\Sigma_\Lambda^R(\nu)}. \quad (29)$$

The corresponding single-scale propagator then reads

$$S^R(\nu) = -\frac{i}{\nu} \frac{[\Theta_\Lambda^R(\nu)]^2 (\nu - \varepsilon_d + i\Delta)}{[\nu - \varepsilon_d + i\Delta - \Theta_\Lambda^R(\nu)\Sigma_\Lambda^R(\nu)]^2}. \quad (30)$$

With this choice, all causality relations and FDRs are satisfied. However, this regulator has two drawbacks compared to the other flow schemes: First, it does not produce a parameter sweep, as  $\Lambda$  does not directly correspond to a physical parameter. Second, computations become ever more challenging for smaller  $\Lambda$ : Even if all correlation functions are reasonably smooth in frequency space for  $\Lambda = 0$ , for small but finite  $\Lambda$ , they exhibit sharp features. While this is not an issue for finite-temperature Matsubara calculations, where only sums over discrete Matsubara frequencies are performed, it turns out to be a major inconvenience in the Keldysh context.

#### IV. CONCLUSION

In this paper, we outline the structure and design of our C++ codebase for diagrammatic calculations of the AM in the Keldysh formalism. We explained the building blocks for representing real-frequency correlation functions and the central routines used to compute them. We elaborated on all performance-critical aspects, allowing one to handle the three-dimensional frequency dependence of the four-point vertex, and summarized the implementation of the parquet and mFRG equations. By discussing the most convenient features of the codebase—modularity, flexibility, performance, and scalability—but also some of its design flaws in detail, we hope to provide guidance and inspiration to others who plan to write code for similar purposes.

Our codebase forms the basis for numerous future projects involving the dynamical correlation functions of electronic many-body systems. Since the AM is very well understood, we want to generalize our treatment to more complicated models with unexplored physics, like lattice models, possibly including multiple bands. The main problem in that regard is the numerical complexity. In addition to their real-frequency Keldysh structure, all functions would acquire momentum dependencies and orbital indices. Parametrizing those naively appears prohibitively costly. Fortunately, the new quantum tensor cross interpolation (QTCI) scheme<sup>76–78</sup> is currently being developed, which can be used to obtain highly compressed tensor network representations of correlation functions and promises exponential reductions in computational costs. It remains to be seen how efficiently the Keldysh four-point vertex can be compressed using this method. If it turned out to be highly compressible, one could combine the diagrammatic approaches outlined here with non-perturbative results from dynamical mean-field theory to access

truly strongly correlated parameter regimes (see related works<sup>79–81</sup> in the MaF). In particular, computing non-local real-frequency dynamical vertex corrections beyond DMFT for observables like optical conductivities with high precision is a formidable long-term goal.

Another possible future direction relates to nonequilibrium phenomena, for example, the influence of the full four-point vertex on observables like differential conductivities.<sup>12,82</sup> Nonequilibrium physics has been the most popular application of the KF in the past, and the AM with a finite bias voltage is tractable with only a minor increase in both the numerical costs and the implementation effort.

In order to leverage ongoing efforts in the QTCI framework, an interface to the corresponding Julia package<sup>83</sup> would be required. Given that, in recent years, multiple Julia codes have been developed to perform calculations of two-particle correlation functions,<sup>84–87</sup> it would be natural to switch to that language in the future, especially since it allows much simpler structures and, in general, performs almost as well as C++.

## ACKNOWLEDGMENTS

We acknowledge Marc Ritter, Marcel Gievers, Benedikt Schneider, Björn Sbierski, Nils Niggemann, Dominik Kiese, Aiman Al-Eryani, Lennart Klebl, and Jacob Beyer for insightful conversations about fRG code structures over the years, Henri Menke for the suggestion to use a combination of Doxygen, Sphinx, and Breathe for the technical code documentation, and Markus Frankenbach for comments on the paper.

N.R. acknowledges funding from a graduate scholarship from the German Academic Scholarship Foundation (“Studienstiftung des deutschen Volkes”) and additional support from the “Marianne-Plehn-Programm” of the state of Bavaria. A.G. and J.v.D. were supported by the Deutsche Forschungsgemeinschaft under Germany’s Excellence Strategy EXC-2111 (Project No. 390814868) and the Munich Quantum Valley, supported by the Bavarian state government with funds from the Hightech Agenda Bayern Plus. This work was supported by Grant No. INST 86/1885-1 FUGG of the German Research Foundation (DFG). F.B.K. acknowledges support from the Alexander von Humboldt Foundation through the Feodor Lynen Fellowship. The Flatiron Institute is a division of the Simons Foundation.

## AUTHOR DECLARATIONS

### Conflict of Interest

The authors have no conflicts to disclose.

### Author Contributions

N.R. and A.G. contributed equally to this work.

**Nepomuk Ritz:** Funding acquisition (equal); Software (supporting); Writing – original draft (lead); Writing – review & editing (lead).  
**Anxiang Ge:** Software (lead); Writing – review & editing (supporting).  
**Elias Walter:** Conceptualization (lead); Software (equal).  
**Santiago Aguirre:** Conceptualization (equal); Software (equal).  
**Jan von Delft:** Funding acquisition (lead); Supervision (equal);

Writing – review & editing (equal). **Fabian B. Kugler:** Supervision (equal); Writing – review & editing (supporting).

## DATA AVAILABILITY

The code, which includes a link to the technical documentation, can be found online (see Ref. 88). It is published under the MIT license.

## REFERENCES

- 1 A. Ge, N. Ritz, E. Walter, S. Aguirre, J. von Delft, and F. B. Kugler, “Real-frequency quantum field theory applied to the single-impurity Anderson model,” *Phys. Rev. B* **109**, 115128 (2024).
- 2 T. Matsubara, “A new approach to quantum-statistical mechanics,” *Prog. Theor. Phys.* **14**, 351–378 (1955).
- 3 A. A. Abrikosov, L. P. Gorkov, and I. E. Dzyaloshinski, *Methods of Quantum Field Theory in Statistical Physics* (Courier Corporation, 2012).
- 4 G. Baym and N. D. Mermin, “Determination of thermodynamic Green’s functions,” *J. Math. Phys.* **2**, 232–234 (1961).
- 5 M. Jarrell and J. Gubernatis, “Bayesian inference and the analytic continuation of imaginary-time quantum Monte Carlo data,” *Phys. Rep.* **269**, 133–195 (1996).
- 6 A. W. Sandvik, “Stochastic method for analytic continuation of quantum Monte Carlo data,” *Phys. Rev. B* **57**, 10287–10290 (1998).
- 7 J. Schött, I. L. M. Locht, E. Lundin, O. Grånäs, O. Eriksson, and I. Di Marco, “Analytic continuation by averaging Padé approximants,” *Phys. Rev. B* **93**, 075104 (2016).
- 8 A. Ge, J. Halbinger, S.-S. B. Lee, J. von Delft, and F. B. Kugler, “Analytic continuation of multipoint correlation functions,” *Ann. Phys.* **536**, 2300504 (2024).
- 9 J. Kroha, P. Wölfle, and T. A. Costi, “Unified description of Fermi and non-Fermi liquid behavior in a conserving slave boson approximation for strongly correlated impurity models,” *Phys. Rev. Lett.* **79**, 261–264 (1997).
- 10 J. Kroha and P. Wölfle, “Fermi and non-Fermi liquid behavior of local moment systems within a conserving slave boson theory,” *Acta Phys. Pol., B* **29**, 3781 (1998).
- 11 S. G. Jakobs, “Functional renormalization group studies of quantum transport through mesoscopic systems,” Ph.D. thesis, RWTH Aachen, 2010.
- 12 S. G. Jakobs, M. Pletyukhov, and H. Schoeller, “Nonequilibrium functional renormalization group with frequency-dependent vertex function: A study of the single-impurity Anderson model,” *Phys. Rev. B* **81**, 195109 (2010).
- 13 S. G. Jakobs, M. Pletyukhov, and H. Schoeller, “Properties of multi-particle Green’s and vertex functions within Keldysh formalism,” *J. Phys. A: Math. Theor.* **43**, 103001 (2010).
- 14 F. B. Kugler, S.-S. B. Lee, and J. von Delft, “Multipoint correlation functions: Spectral representation and numerical evaluation,” *Phys. Rev. X* **11**, 041006 (2021).
- 15 S.-S. B. Lee, F. B. Kugler, and J. von Delft, “Computing local multipoint correlators using the numerical renormalization group,” *Phys. Rev. X* **11**, 041007 (2021).
- 16 F. B. Kugler and J. von Delft, “Multiloop functional renormalization group for general models,” *Phys. Rev. B* **97**, 035162 (2018).
- 17 F. B. Kugler and J. von Delft, “Multiloop functional renormalization group that sums up all parquet diagrams,” *Phys. Rev. Lett.* **120**, 057403 (2018).
- 18 F. B. Kugler and J. von Delft, “Derivation of exact flow equations from the self-consistent parquet relations,” *New J. Phys.* **20**, 123029 (2018).
- 19 J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, “Julia: A fresh approach to numerical computing,” *SIAM Rev.* **59**, 65–98 (2017).
- 20 N. D. Matsakis and F. S. Klock II, “The rust language,” in *ACM SIGAda Ada Letters* (ACM, 2014), Vol. 34, pp. 103–104.
- 21 See <https://www.modular.com/max/mojo> for Mojo—The programming language for all AI developers.
- 22 B. Stroustrup, in *Programming—Principles and Practice Using C++*, 3rd ed. (Addison Wesley, Boston, MA, 2023).

- <sup>23</sup>R. Bulla, T. A. Costi, and T. Pruschke, "Numerical renormalization group method for quantum impurity systems," *Rev. Mod. Phys.* **80**, 395–450 (2008).
- <sup>24</sup>V. Zlatić and B. Horvatić, "Series expansion for the symmetric Anderson Hamiltonian," *Phys. Rev. B* **28**, 6904–6906 (1983).
- <sup>25</sup>P. B. Wiegmann and A. M. Tsvelick, "Exact solution of the Anderson model: I," *J. Phys. C: Solid State Phys.* **16**, 2281 (1983).
- <sup>26</sup>P. W. Anderson, "Localized magnetic states in metals," *Phys. Rev.* **124**, 41–53 (1961).
- <sup>27</sup>A. Kamenev, in *Field Theory of Non-Equilibrium Systems*, 2nd ed. (Cambridge University Press, 2023).
- <sup>28</sup>E. Walter, "Real-frequency dynamics of quantum impurity models studied with fRG, NRG, CFT," Ph.D. thesis, LMU München, 2021.
- <sup>29</sup>L. V. Keldysh, "Diagram technique for nonequilibrium processes," *Sov. Phys. JETP* **20**, 1018–1026 (1965).
- <sup>30</sup>J. Schwinger, "Brownian motion of a quantum oscillator," *J. Math. Phys.* **2**, 407–432 (1961).
- <sup>31</sup>L. P. Kadanoff and G. A. Baym, *Quantum Statistical Mechanics* (Benjamin, New York, 1962).
- <sup>32</sup>ISO/IEC 14882:2017 - Programming languages—C++, International Organization for Standardization, 2017.
- <sup>33</sup>Kitware, Inc., CMake documentation, <https://cmake.org/documentation/>.
- <sup>34</sup>GNU Scientific Library Reference Manual: For GSL version 1.12, 3rd ed., edited by M. Galassi (Network Theory, Bristol, 2009).
- <sup>35</sup>See <https://www.boost.org/> for Boost C++ Libraries.
- <sup>36</sup>G. Guennebaud, B. Jacob *et al.*, Eigen: A C++ linear algebra library, <http://eigen.tuxfamily.org>.
- <sup>37</sup>The HDF Group, Hierarchical data format, version 5, <https://www.hdfgroup.org/HDF5/>.
- <sup>38</sup>See <https://www.openmp.org/> for the OpenMP API specification for parallel programming.
- <sup>39</sup>See <https://hpc.nmsu.edu/discovery/mpi/introduction/> for message passing interface: High performance computing.
- <sup>40</sup>D. van Heesch, Doxygen: A documentation system for C++, C, Java, Python and other languages, <https://www.doxygen.nl/>.
- <sup>41</sup>See <https://www.sphinx-doc.org/> for Sphinx Documentation Generator.
- <sup>42</sup>See <https://breathe.readthedocs.io/en/latest/> for Breathe "latest" documentation.
- <sup>43</sup>G. Rohringer, H. Hafermann, A. Toschi, A. A. Katanin, A. E. Antipov, M. I. Katsnelson, A. I. Lichtenstein, A. N. Rubtsov, and K. Held, "Diagrammatic routes to nonlocal correlations beyond dynamical mean field theory," *Rev. Mod. Phys.* **90**, 025003 (2018).
- <sup>44</sup>A. Georges, G. Kotliar, W. Krauth, and M. J. Rozenberg, "Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions," *Rev. Mod. Phys.* **68**, 13–125 (1996).
- <sup>45</sup>A. Toschi, A. A. Katanin, and K. Held, "Dynamical vertex approximation: A step beyond dynamical mean-field theory," *Phys. Rev. B* **75**, 045118 (2007).
- <sup>46</sup>C. Taranto, S. Andergassen, J. Bauer, K. Held, A. Katanin, W. Metzner, G. Rohringer, and A. Toschi, "From infinite to two dimensions through the functional renormalization group," *Phys. Rev. Lett.* **112**, 196402 (2014).
- <sup>47</sup>N. Wentzell, G. Li, A. Tagliavini, C. Taranto, G. Rohringer, K. Held, A. Toschi, and S. Andergassen, "High-frequency asymptotics of the vertex function: Diagrammatic parametrization and algorithmic implementation," *Phys. Rev. B* **102**, 085106 (2020).
- <sup>48</sup>A. A. Katanin, "Fulfillment of Ward identities in the functional renormalization group approach," *Phys. Rev. B* **70**, 115109 (2004).
- <sup>49</sup>F. B. Kugler, "Improved estimator for numerical renormalization group calculations of the self-energy," *Phys. Rev. B* **105**, 245132 (2022).
- <sup>50</sup>E. Wang and U. Heinz, "Generalized fluctuation-dissipation theorem for nonlinear response functions," *Phys. Rev. D* **66**, 025008 (2002).
- <sup>51</sup>C. Hille, F. B. Kugler, C. J. Eckhardt, Y.-Y. He, A. Kauch, C. Honerkamp, A. Toschi, and S. Andergassen, "Quantitative functional renormalization group description of the two-dimensional Hubbard model," *Phys. Rev. Res.* **2**, 033372 (2020).
- <sup>52</sup>D. Rohe, "Hierarchical parallelisation of functional renormalisation group calculations—hp-fRG," *Comput. Phys. Commun.* **207**, 160–169 (2016).
- <sup>53</sup>W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, in *Numerical Recipes—The Art of Scientific Computing*, 3rd ed. (Cambridge University Press, 2007).
- <sup>54</sup>A modern, C++-native, test framework for unit-tests, TDD and BDD.
- <sup>55</sup>D. C. Langreth, "Friedel sum rule for Anderson's model of localized impurity states," *Phys. Rev.* **150**, 516–518 (1966).
- <sup>56</sup>N. E. Bickers, "Self-consistent many-body theory for condensed matter systems," in *Theoretical Methods for Strongly Correlated Electrons, CRM Series in Mathematical Physics*, edited by D. Sénéchal, A.-M. Tremblay, and C. Bourbonnais (Springer, New York, 2004).
- <sup>57</sup>D. G. Anderson, "Iterative procedures for nonlinear integral equations," *J. ACM* **12**, 547–560 (1965).
- <sup>58</sup>H. F. Walker and P. Ni, "Anderson acceleration for fixed-point iterations," *SIAM J. Numer. Anal.* **49**, 1715–1735 (2011).
- <sup>59</sup>W. Metzner, M. Salmhofer, C. Honerkamp, V. Meden, and K. Schönhammer, "Functional renormalization group approach to correlated fermion systems," *Rev. Mod. Phys.* **84**, 299–352 (2012).
- <sup>60</sup>C. Wetterich, "Exact evolution equation for the effective potential," *Phys. Lett. B* **301**, 90–94 (1993).
- <sup>61</sup>F. B. Kugler and J. von Delft, "Fermi-edge singularity and the functional renormalization group," *J. Phys.: Condens. Matter* **30**, 195501 (2018).
- <sup>62</sup>F. Krien, "Efficient evaluation of the polarization function in dynamical mean-field theory," *Phys. Rev. B* **99**, 235106 (2019).
- <sup>63</sup>F. Krien, A. Valli, and M. Capone, "Single-boson exchange decomposition of the vertex function," *Phys. Rev. B* **100**, 155149 (2019).
- <sup>64</sup>F. Krien and A. Valli, "Parquetlike equations for the Hedin three-leg vertex," *Phys. Rev. B* **100**, 245147 (2019).
- <sup>65</sup>F. Krien, A. I. Lichtenstein, and G. Rohringer, "Fluctuation diagnostic of the nodal/antinodal dichotomy in the Hubbard model at weak coupling: A parquet dual fermion approach," *Phys. Rev. B* **102**, 235133 (2020).
- <sup>66</sup>F. Krien, A. Valli, P. Chalupa, M. Capone, A. I. Lichtenstein, and A. Toschi, "Boson-exchange parquet solver for dual fermions," *Phys. Rev. B* **102**, 195131 (2020).
- <sup>67</sup>F. Krien, A. Kauch, and K. Held, "Tiling with triangles: Parquet and GWy methods unified," *Phys. Rev. Res.* **3**, 013149 (2021).
- <sup>68</sup>M. Gievers, E. Walter, A. Ge, J. von Delft, and F. B. Kugler, "Multiloop flow equations for single-boson exchange fRG," *Eur. Phys. J. B* **95**, 108 (2022).
- <sup>69</sup>K. Fraboulet, S. Heinzelmann, P. M. Bonetti, A. Al-Eryani, D. Vilardi, A. Toschi, and S. Andergassen, "Single-boson exchange functional renormalization group application to the two-dimensional Hubbard model at weak coupling," *Eur. Phys. J. B* **95**, 202 (2022).
- <sup>70</sup>J. R. Cash and A. H. Karp, "A variable order Runge-Kutta method for initial value problems with rapidly varying right-hand sides," *ACM Trans. Math. Software* **16**, 201–222 (1990).
- <sup>71</sup>C. Honerkamp, D. Rohe, S. Andergassen, and T. Enss, "Interaction flow method for many-fermion systems," *Phys. Rev. B* **70**, 235115 (2004).
- <sup>72</sup>C. Honerkamp and M. Salmhofer, "Temperature-flow renormalization group and the competition between superconductivity and ferromagnetism," *Phys. Rev. B* **64**, 184516 (2001).
- <sup>73</sup>B. Schneider, J. Reuther, M. G. Gonzalez, B. Sbierski, and N. Niggemann, "Temperature flow in pseudo-Majorana functional renormalization for quantum spins," *Phys. Rev. B* **109**, 195109 (2024).
- <sup>74</sup>A. Tagliavini, C. Hille, F. B. Kugler, S. Andergassen, A. Toschi, and C. Honerkamp, "Multiloop functional renormalization group for the two-dimensional Hubbard model: Loop convergence of the response functions," *SciPost Phys.* **6**, 009 (2019).
- <sup>75</sup>P. Chalupa-Gantner, F. B. Kugler, C. Hille, J. von Delft, S. Andergassen, and A. Toschi, "Fulfillment of sum rules and Ward identities in the multiloop functional renormalization group solution of the Anderson impurity model," *Phys. Rev. Res.* **4**, 023050 (2022).
- <sup>76</sup>H. Shinaoka, M. Wallerberger, Y. Murakami, K. Nogaki, R. Sakurai, P. Werner, and A. Kauch, "Multiscale space-time ansatz for correlation functions of quantum systems based on quantum tensor trains," *Phys. Rev. X* **13**, 021015 (2023).

- <sup>77</sup>Y. Núñez Fernández, M. Jeannin, P. T. Dumitrescu, T. Kloss, J. Kaye, O. Parcollet, and X. Waintal, “Learning Feynman diagrams with tensor trains,” *Phys. Rev. X* **12**, 041018 (2022).
- <sup>78</sup>M. K. Ritter, Y. Núñez Fernández, M. Wallerberger, J. von Delft, H. Shinaoka, and X. Waintal, “Quantics tensor cross interpolation for high-resolution parsimonious representations of multivariate functions,” *Phys. Rev. Lett.* **132**, 056501 (2024).
- <sup>79</sup>M. Kitatani, T. Schäfer, H. Aoki, and K. Held, “Why the critical temperature of high- $T_c$  cuprate superconductors is so low: The importance of the dynamical vertex structure,” *Phys. Rev. B* **99**, 041115 (2019).
- <sup>80</sup>D. Vilardi, C. Taranto, and W. Metzner, “Antiferromagnetic and  $d$ -wave pairing correlations in the strongly interacting two-dimensional Hubbard model from the functional renormalization group,” *Phys. Rev. B* **99**, 104501 (2019).
- <sup>81</sup>P. M. Bonetti, A. Toschi, C. Hille, S. Andergassen, and D. Vilardi, “Single-boson exchange representation of the functional renormalization group for strongly interacting many-electron systems,” *Phys. Rev. Res.* **4**, 013034 (2022).
- <sup>82</sup>T. Fujii and K. Ueda, “Perturbative approach to the nonequilibrium Kondo effect in a quantum dot,” *Phys. Rev. B* **68**, 155310 (2003).
- <sup>83</sup>M. K. Ritter *et al.*, *Quanticsctci.jl*, 2022.
- <sup>84</sup>D. Kiese, T. Müller, Y. Iqbal, R. Thomale, and S. Trebst, “Multiloop functional renormalization group approach to quantum spin systems,” *Phys. Rev. Res.* **4**, 023185 (2022).
- <sup>85</sup>D. Kiese, A. Ge, N. Ritz, J. von Delft, and N. Wentzell, “MatsubaraFunctions.jl: An equilibrium Green’s function library in the Julia programming language,” *SciPost Phys. Codebases* **24** (2024).
- <sup>86</sup>D. Kiese, A. Ge, N. Ritz, J. von Delft, and N. Wentzell, “Codebase release 0.1 for MatsubaraFunctions.jl,” *SciPost Phys. Codebases* **24-r0.1** (2024).
- <sup>87</sup>N. Niggemann, *Nilsniggemann/pmfrg.jl: v2.1.9*, 2023.
- <sup>88</sup>N. Ritz and A. Ge, *KeldyshQFT: A C++ Codebase for real-frequency multiloop functional renormalization group and parquet computations for the single-impurity Anderson model*.