UvA-DARE (Digital Academic Repository)

Correlation functions of integrable models: A description of the ABACUS algorithm

Caux, J.S.

DOI

[10.1063/1.3216474](https://doi.org/10.1063/1.3216474)

Publication date

2009

Document Version

Final published version

Published in

Journal of Mathematical Physics

[Link to publication](#)

Citation for published version (APA):

Caux, J. S. (2009). Correlation functions of integrable models: A description of the ABACUS algorithm. *Journal of Mathematical Physics*, 50(9), 095214. <https://doi.org/10.1063/1.3216474>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

Correlation functions of integrable models: A description of the ABACUS algorithm

Jean-Sébastien Caux^{a)}

Institute for Theoretical Physics, Universiteit van Amsterdam, Valckenierstraat 65, 1018 XE Amsterdam, The Netherlands

(Received 12 May 2009; accepted 7 August 2009; published online 24 September 2009)

Recent developments in the theory of integrable models have provided the means of calculating dynamical correlation functions of some important observables in systems such as Heisenberg spin chains and one-dimensional atomic gases. This article explicitly describes how such calculations are generally implemented in the ABACUS C++ library, emphasizing the universality in treatment of different cases coming as a consequence of unifying features within the Bethe ansatz. © 2009 American Institute of Physics. [DOI: [10.1063/1.3216474](https://doi.org/10.1063/1.3216474)]

I. INTRODUCTION

The Bethe ansatz,¹ introduced in 1931 (very shortly after the appearance of quantum mechanics), is often viewed as an arcane subject, manifestly interesting for mathematically minded theorists but not for a broad audience. At least two good reasons for this can be given. First, the subject of integrability is rather specialized and demanding in itself;^{2–5} it is after all made up of a collection of somewhat elaborate methods having only limited general applicability, and is therefore not typically considered worthy of inclusion in the standard condensed matter curriculum (the fact remains that the technology developed around integrable models has proven its usefulness in many ways, perhaps most importantly not only in providing many contributions to our understanding of strongly correlated physics in one dimension⁶ but also in providing reliable beacons for the testing and benchmarking of more widely used field-theory or numerics-based methods for such systems). Second, and perhaps more importantly, is that the experimental relevance of integrable systems remains limited. Good realizations of one-dimensional integrable quantum systems do exist, but are the exception rather than the rule; when a good realization is found, integrability often cannot provide quantitative predictions for the most important observable quantities.

This last difficulty is related to the long-standing inability of the Bethe ansatz to provide results for even simple correlation functions of local physical observables. The source of this difficulty lies in the admittedly unwieldy form of Bethe wave functions, which are composed of a factorially large sum of free waves with nontrivial relative amplitudes. The action of local operators on these wave functions is difficult to write down, and since this forms the starting point in computing any correlation function, progress has been tedious, slow, and limited.

Recent years have, however, seen that some important developments occur at a rapid pace in the calculation of dynamical correlations for finite systems at zero temperature. This stems from the appearance of a number of important results in the field of integrability. First and foremost is Slavnov's theorem on scalar products, providing initial results on form factors and correlation functions.^{7,8} The general problem of describing the action of local operators on Bethe wave functions was then resolved with the breakthrough solution of the quantum inverse problem for spin chains,^{9,10} opening the door to the computation of general correlation functions of the XXZ chain. The problem of obtaining correlation functions from these representations was then initiated in Refs. 11–13 with a study of traces over two-particle intermediate states. The ABACUS algorithm

^{a)}Electronic mail: j.s.caux@uva.nl.

was thereafter developed in order to tackle the summation of all important general multiparticle excitations in Heisenberg spin chains and the one-dimensional Bose gas.^{14–17} The present article only provides a “conceptual-level” description of the features and functioning of ABACUS, only emphasizing unifying features of this class of calculations without going deeply into specific examples. My main objective in this article is to give a more or less detailed discussion of the strategies adopted in ABACUS, to define some terminology which will prove useful in future works, and to make some observations on the results obtained. A separate user guide¹⁸ will provide implementation-level details of the library and its usage, including important primary class and function declarations. Extensive and systematic results for various models will form the subject of specialized publications. The subject of integrability itself will be covered in a separate set of lecture notes.¹⁹ The general field of correlation functions of strongly correlated one-dimensional quantum models is the subject of an upcoming review.²⁰

II. DEFINING THE PROBLEM

Let us for definiteness consider a quantum-mechanical model on a lattice, with Hamiltonian H which we assume for convenience (and without true loss of generality) to act in a Hilbert space formed by the tensoring of finitely many single-site local Hilbert spaces. Letting N be the number of such sites and using $j=1, \dots, N$ as a site label, all operators can be reconstructed from the self-dual set $\mathcal{O}_j^a$ of local physical operators, where a is some on-site representation index.

The general problem which we will address is the calculation of two-point zero-temperature equilibrium correlation functions of the form

$$\langle \mathcal{O}_j^a(t) \mathcal{O}_{j'}^{\bar{a}}(0) \rangle, \quad (1)$$

where $\langle \dots \rangle$ represents the ground-state expectation value and $\mathcal{O}_j^{\bar{a}} = (\mathcal{O}_j^a)^\dagger$ for arbitrary time differences t and distances $j-j'$. Equivalently, we can consider the Fourier transform of this correlation function, which (by a slight abuse of terminology) will be called the dynamical structure factor (DSF),

$$S^{a\bar{a}}(k, \omega) = \frac{1}{N} \sum_{j,j'=1}^N e^{-ik(j-j')} \int_{-\infty}^{\infty} dt e^{i\omega t} \langle \mathcal{O}_j^a(t) \mathcal{O}_{j'}^{\bar{a}}(0) \rangle. \quad (2)$$

A complete knowledge of the DSF for all momenta and frequencies naturally allows to reconstruct (1) without approximation. In reality, however, (2) itself is of more direct usefulness since it is often the quantity which can be related to energy- and momentum-resolved experimental measurements in linear response.

Calculating the DSF for interacting models (by which is meant models where multiparticle states are not created by simple products of single-particle creation operators) is a very complex task. The product of quantum operators acting at different times and places is best dealt with by introducing a summation over intermediate eigenstates (labeled by Greek indices), allowing to resolve the time dependence explicitly. Using the space Fourier transform $\mathcal{O}_k^a = \sum_j e^{-ikj} \mathcal{O}_j^a$, summing over lattice sites, and performing the time integration lead to the Lehmann series representation,

$$S^{a\bar{a}}(k, \omega) = \frac{2\pi}{N} \sum_{\mu} |\langle \lambda^0 | \mathcal{O}_k^a | \mu \rangle|^2 \delta(\omega - E_{\mu} + E_0), \quad (3)$$

where E_0 is the energy of the ground state $|\lambda^0\rangle$. The matrix elements $\langle \lambda | \mathcal{O} | \mu \rangle$ of local operators in the eigenstates basis are also known as form factors (for continuum models); their norm square is also often called transition rate in the literature, in view of the correspondence of (3) to Fermi's golden rule.

We can make a simple “wish list” of the elements needed to provide a computation of such DSFs:

- (1) an orthonormal eigenstate basis;
- (2) form factors (matrix elements) $\langle \lambda | \mathcal{O}_k^a | \mu \rangle$ of operators $\mathcal{O}_k^a$ in this basis;
- (3) a way to perform the summation over intermediate states.

In general, this is an immensely complicated problem, which can typically be performed for models which correspond or can be mapped to free particles, or which benefit from a conformal field theory description. Simplistic treatments of interacting systems, in general, do not provide us with any of the three needed elements.

In one dimension, however, the technology of integrability does provide us with some of the three elements above, for a number of interesting models. The Bethe ansatz provides the eigenstates basis, and the algebraic Bethe ansatz (together with the solution of the quantum inverse problem) provides us with form factors. The summation over intermediate states remains, however, an open problem, for which partial analytical results are possible in some restricted circumstances, but which still demands more often than not a numerical solution.

The ABACUS library has recently been developed in order to perform the calculation of DSFs of the known integrable models for which eigenstates and form factors can be explicitly constructed. ABACUS (Algebraic Bethe Ansatz-based Computation of Universal Structure factors) not only provides functions for representing and constructing eigenstates, calculating form factors of a number of physical operators of interest, but also implements a scanning algorithm over intermediate states allowing to construct the correlations explicitly. My main motivations in developing ABACUS was primarily not only to provide quantitative predictions first for experiments in low-dimensional magnetism and later on cold atomic gases but also to provide results allowing cross-checking of more generally applicable alternative theoretical (analytical or numerical) approaches to the dynamics of strongly correlated systems.

The method which is presented here, and the results that it provides, undoubtedly suffer from a number of disadvantages and shortcomings. Among those, we can mention the following.

- (1) It is restricted to one-dimensional, Bethe ansatz integrable models.
- (2) It is restricted to the simplest known models, excluding all “nested” systems.
- (3) It only applies to zero-temperature DSFs.
- (4) It can only treat finite systems.
- (5) It does not provide an exact, closed-form analytical expression, but only a numerical result.

Restriction 1 is insurmountable: the whole ABACUS edifice begins and ends with the Bethe ansatz. At best, what will be achievable is an extension of the framework to provide a description of correlations in perturbed integrable models. Restriction 2 is probably temporary and exists simply because it is not yet known whether economical expressions for form factors of local operators in nested systems exist or can be realistically found. It is also questionable whether restriction 3 is temporary or not: on the one hand, Bethe ansatz can deal with finite temperatures quite straightforwardly for equilibrium expectation values. On the other hand, the problem of finite temperature Gibbs traces opens up a Pandora’s box of complications within the ABACUS logic, and it is not yet clear whether these can be overcome. Restriction 4 comes from the fact that the eigenstates dealt with must be given by finite numbers of rapidities, and that the Hilbert space must have a measure of finiteness so that summing over the important intermediate states can be done in a finite time. It is thus not possible for ABACUS to give results in the strict infinite-size limit. The size dependence of the results is a complex affair, discussed in a later section. Restriction 5 is the most severe. ABACUS is and will remain numerical; whether simple exact closed-form expressions even exist for general DSFs is itself more than debatable at present. There are very good mathematical reasons to believe that such expressions cannot be found, except in various simplifying circumstances and/or limits.

For completeness, and in order to relieve the discouragement of the reader, let us now list the advantages of ABACUS.

- (1) It works for very important “canonical” strongly correlated models such as Heisenberg chains and the Lieb–Liniger model.

- (2) Its implementation is relatively universal to the set of integrable models.
- (3) It provides extremely accurate results for large systems.
- (4) Its accuracy is to a large degree energy and momentum independent.

Advantage 1 is not strictly an advantage, but is worth emphasizing anyway: the models treated are strongly interacting systems which *do* have experimental applicability and an important role as beacons of reference for refining and fine-tuning alternative approaches. Advantage 2 comes from the fact that in the field of integrability, if one knows how to perform a “trick” in one model, one can often apply the same trick to most if not all other integrable models. ABACUS is built from the start with the intension of exploiting this “universality” within the Bethe ansatz. Advantage 3 is the most prominent: the accuracy level can be assessed with sum rules, and saturation levels beyond 99% are achievable for systems with many hundreds of particles, even in the more difficult limits. Relaxing this requirement ever so slightly, say to 90%, one can go to systems with thousands of particles. In any case, the system sizes achievable are comfortably much, much beyond anything which will ever be realistically achievable using exact diagonalization. Moreover, the system sizes attainable are also high enough that finite-size effects are drastically reduced, except perhaps in the vicinity of excitation thresholds, where the correspondence to singularities which develop in the infinite-size limit seems difficult to obtain. The fact that large systems can be treated is a consequence of the amount of preliminary information which is provided by integrability in the first place, and which can be built directly into the algorithm. The fact that ABACUS can give extremely accurate results for finite systems can itself be seen as an advantage, when applications to nano/mesoscopic quantum systems are considered. Advantage 4 is an interesting one, since many other methods (such as bosonization and conformal field theory) are only valid at low energies. The common lore is that only “universal” low-energy features should be computed by theorists, since these are the only ones which are unambiguously observable in experiments. A more careful assessment, however, teaches that such low-energy limits are precisely *not* accessible experimentally, and that reliable data readout has to be done at higher energy scales. A less misleading statement about “universality” is thus that the extrapolation via scaling of the low-energy, universal result (which is not itself consistently observable) to somewhat higher energy scales (which are) can provide a good fit. However, the limitations of “universal” approaches quickly becomes severe: in the case of spin chains, for example, inelastic neutron scattering provides data over the full Brillouin zone, showing clearly the band curvature and interaction effects which cannot be quantified using bosonization. ABACUS shows its strength here by being able to quantify these effects: it can compute the DSF at any energy scale, giving very precise data for any region of the momentum/energy plane. Since the Bethe ansatz used by ABACUS gives exact wave functions irrespective of their energy, the accuracy of the data becomes more or less energy independent (since more types of excited states can live at higher energies, the accuracy there becomes limited by the restricted number of higher excited states that can be constructed; there thus remains an energy dependence to the accuracy, but it is weak). This feature of ABACUS makes it, at least in the mind of the author, the most appropriate method to give *quantitative* predictions for experiments.

Let us now get down to the specifics. The bulk of the paper is organized as follows, mirroring the three elements of the above “wish list.” Section III provides an extensive discussion of the conventions and terminology used in the classification and representation of eigenstates within ABACUS. Section IV then discusses how individual eigenstates and their norm are obtained, and how form factors of local operators are computed. Section V describes how the Hilbert space of intermediate states is scanned in order to achieve optimal results for the DSF. Section VI provides some example results obtained from ABACUS, with a discussion of some of their specific features. Finally, Sec. VII presents a discussion of some generic features of the approach, and the paper ends with conclusions and perspectives for the future.

III. CLASSIFICATION AND REPRESENTATION OF EIGENSTATES

This section aims at providing a representation of eigenstates of Bethe ansatz integrable models which is tailor made for the purposes of calculating correlation functions. Such a repre-

sentation will rely only on very generic features of Bethe eigenstates, so the discussion will not mention any model, in particular (except to highlight exceptional aspects of specific cases).

Our integrable model H will have a Fock space $\mathcal{F}$ of dimensionality $\dim(\mathcal{F})$, which can be finite (for e.g., finite lattice models) or infinite (for e.g., continuum models with no UV cutoff). Typically, we will be able to easily separate this Fock space into a tensor product of spaces $\mathcal{F}_M$ with fixed number M of “particles,” with M being the charge of one of the trivially conserved quantities.²¹

A. String basis: Configurations

The Bethe ansatz gives us explicit wave functions for eigenstates of H in a fixed-charge subspace $\mathcal{F}_M$, parametrized in terms of M rapidities λ_a , $a=1, \dots, M$. In all generality, the rapidities of a given Bethe ansatz solvable model live in the complex plane, and this introduces immense complications in the attempt to prove that the basis of Bethe eigenstates is complete. Under broad and reasonable assumptions, however, the common lore is that this basis *is* complete,¹ and that states can be understood and classified in terms of collections of (possibly deformed) generalized “string” patterns.²² For a given model, let us call N_s the total number of possible string patterns. We let the label j run through the values $1, \dots, N_s$, with increasing j representing strings of increasing charge. In what follows, we call j the string *level*. A given eigenstate will have specific numbers M_j of strings of level j . We call a set $\{M_j\}$, $j=1, \dots, N_s$ a *base*. The total charge $M = \sum_j l_j M_j$, where l_j is the charge of a string of level j , is simply called the *charge* of a base, and we sometimes write the charge of a base as a subscript (i.e., $\{M_j\}_M$). The total number of strings, $M_{st} = \sum_j M_j \leq M$ is called the *string charge* of a base.

Since two-particle scattering preserves strings, the space $\mathcal{F}_M$ further subdivides into subspaces $\mathcal{F}_{\{M_j\}_M}$ having fixed-charge- M bases. In other words, each partitioning of M into a charge- M base $\{M_j\}_M$ generates an independent Hilbert space.

The exponential form of the Bethe equations is a set of transcendental equations for the rapidities. When written in logarithmic form, these equations can be interpreted as a mapping between the set of rapidities $\{\lambda_\alpha^j\}$ and a set of quantum numbers $\{I_\alpha^j\}$ with $j=1, \dots, N_s$ and $\alpha=1, \dots, M_j$ at each level,

$$N\theta_{\text{kin}}^j(\lambda_\alpha^j) + \sum_{k=1}^{N_s} \sum_{\alpha=1}^{M_k} \theta_{\text{scat}}^{jk}(\lambda_\alpha^j, \lambda_\beta^k) = 2\pi I_\alpha^j, \quad (4)$$

where θ_{kin}^j is a kinetic function and $\theta_{\text{scat}}^{jk}$ is a scattering phase shift kernel. Since these equations make use of the string hypothesis (of perfectly undeformed strings whose deviations are assumed exponentially small in system size), thus reducing the number of unknown parameters from the charge to the string charge, it seems historically appropriate^{23–25} to call them the Bethe–Gaudin–Takahashi (BGT) equations, reserving the name Bethe equations to the original full set (for a number of rapidity parameters equal to the charge) of coupled equations without any string hypothesis made. Although the rapidities are all interdependent, the quantum numbers are independent of one another, modulo applying a simple generalized Pauli principle stating that no two quantum numbers at the same level are allowed to coincide, i.e., for $\alpha \neq \beta \in [1, M_j]$, we must have $I_\alpha^j \neq I_\beta^j$, $\forall j$. The quantum numbers thus provide us with the appropriate labeling of the eigenstates. We thus write $|\{I\}\rangle$ or $|\{\lambda\}\rangle$ for our eigenstates, with the understanding that $\{\lambda\}$ are obtained from solving the BGT equations with the choice $\{I\}$.

The dimensionality of a given $\mathcal{F}_{\{M_j\}_M}$ subspace coincides with the number of allowable choices of quantum numbers, since these obey a generalized Pauli principle.²⁶ These can be counted from the observation that choosing a base fixes the sets of possible quantum numbers $\{I_\alpha^j\}$ which can be made at each level. In other words, the choice of a base $\{M_j\}$ unambiguously defines a set $\{I_\alpha^j\}$, $j=1, \dots, N_s$ of left- and right-limiting quantum numbers $I_{\infty,-}^j$ and $I_{\infty,+}^j$ at each level: any eigenstate in this base must have quantum numbers individually obeying the limiting inequalities $I_\alpha^j \in [-I_{\infty,-}^j, I_{\infty,+}^j]$, $\alpha=1, \dots, M_j$, $j=1, \dots, N_s$.²⁷ The dimensionality of this subspace is thus

FIG. 1. Representation of quantum numbers at level $j=1$. In (a), the set of possible quantum numbers at this level for a given base is shown as open circles. In this case, we must fill in $M_1=5$ of the possibilities defined by the limiting quantum numbers $I_{\infty,-}^1=5$ and $I_{\infty,+}^1=6$. (b) shows the lowest-energy configuration. (c) shows a relatively low-energy excited configuration with two particle-hole pairs (four excitations) at this level, with the particles and holes staying close to the Fermi boundary. Finally, (d) gives an example of a high-energy configuration with three particle-hole pairs (six excitations), with holes going further in and particles further out away from the Fermi boundary.

$$\dim(\mathcal{F}_{\{M_j\}_M}) = \prod_{j=1}^{N_s} \binom{I_{\infty,+}^j + I_{\infty,-}^j + 1}{M_j}. \quad (5)$$

We call a set of quantum numbers $\{\vec{I}_\alpha^j\}$ fulfilling the generalized Pauli principle and the limiting inequalities a *configuration*. A unique quantum number configuration is thus associated with each individual eigenstate via the BGT equations. One subtlety is that not all configurations lead to sensible string structures. If a set of quantum numbers $\{\vec{I}_\alpha^j\}$ is parity invariant, then its associated rapidities will also be parity invariant. If quantum numbers $\vec{I}_\alpha^{j_1}$ and $\vec{I}_\beta^{j_2}$ for different levels j_1 and j_2 within such a set are then simultaneously zero, both string centers $\lambda_\alpha^{j_1}$ and $\lambda_\beta^{j_2}$ vanish according to the BGT equations. If the string charges are such that $l_{j_1} = l_{j_2} \bmod 2$, then the string hypothesis would predict coinciding rapidities modulo exponentially small deviations (e.g., an overlapping three- and five-string on zero rapidity leads to three pairwise superpositions). Since the string hypothesis patently fails in such cases, we call these states *inadmissible*, the usual states being *admissible*. Inadmissible states do give proper wave functions involving (strongly) deviated strings, but the original Bethe equations must be solved. ABACUS ignores inadmissible states in the current implementation, since the numbers of inadmissible states necessarily are suppressed in a $1/N$ fashion as compared to admissible ones, and in any case their contribution to correlation functions can in some cases be shown to vanish identically.²⁸

For the calculation of zero-temperature correlation functions as performed by ABACUS, the most important eigenstates are those for which $M_1 \leq M$ and $M_j \sim O(1)$ for $j > 1$, in other words eigenstates for which only a handful of “higher” strings are present, and most rapidities sit at level $j=1$.

1. Representation of configurations in ABACUS

Let us now describe how a configuration is represented and labeled in ABACUS. Let us start by considering level $j=1$, the level at which we will make use of the most complicated representation. The M_1 quantum numbers of this level are to be chosen within the set of quantum numbers $-I_{\infty,-}^1, -I_{\infty,-}^1 + 1, \dots, I_{\infty,+}^1 - 1, I_{\infty,+}^1$. We represent the possible choices by open circles, as per (a) in Fig. 1.²⁹ In the same figure, (b) represents the lowest-energy configuration, while (c) and (d) represent higher-energy configurations. The fact that (b) represents the lowest-energy configura-

FIG. 2. The four sectors of excitations at level $j=1$. Sectors 0 and 1 are, respectively, the right- and left-hand sides of the base Fermi set, while 2 and 3 are the right and left outsides of the base Fermi set. The excitation type for the upper configuration is 0000. The excitation type for the lower configuration is 1212 (see main text).

tion comes by design of the rapidity parametrization and can be ensured at all levels, for all models.

We often use the term “particle” to denote a given occupied quantum number, representing a string in the eigenstate. In most cases, we call “excitations” of a configuration the holes within the base Fermi set (this being defined as the lowest-energy configuration for a given base), and the particles outside of it. Thus, configurations (b)–(d) in Fig. 1 each have five particles, but example (b) has 0 excitations, while (c) has four, and (d) has six (this is not a strict rule: a simple exception are the two-spinon states of Heisenberg chains in zero field, where the spinons are dispersive “holes,”³⁰ and the “particles” living out of the base Fermi set are fixed (only one quantum number slot is available); since the positions of the two holes completely specify the state, we then say that such states contain two excitations). In circumstances where we need to be more specific, we will distinguish between “dispersive” and “nondispersive” excitations. The latter have only one possible choice for their quantum number, and therefore are not part of an excitation family; the former, on the other hand, have many quantum number possibilities, and in most cases become true excitation branches in the infinite-size limit.

Sectors

Excitations formed by a number of particle-hole pairs can be further separated into different classes, depending on which sides of the base Fermi set they live on. We first define four *sectors* as illustrated in Fig. 2. Sector 0 contains the quantum numbers obeying $I \geq 0$, $I < I_F \equiv M_1/2$ (right-hand side of the base Fermi set, including zero). Sector 1 contains those obeying $I < 0$, $I > -I_F$ (left-hand side of base Fermi set). Sector 2 contains the quantum numbers for which $I > I_F$ and $I \leq I_{\infty,+}^1$, while sector 3 contains those obeying $I < -I_F$ and $I \geq -I_{\infty,-}^1$. The numbers of excitations e_i^1 , $i=0, \dots, 3$ at level $j=1$, combined and written out as $e_3^1 e_2^1 e_1^1 e_0^1$ is called the *level 1 type*.³¹ The fact that excitations come from particle-hole pairs means that at level $j=1$, we have the identity $e_0^1 + e_1^1 = e_2^1 + e_3^1$. Despite this slight redundancy, we still label the type as described above, since it then allows to visualize the overall quantum number excitation pattern of a state with a quick glance. Thus, in Fig. 1, the configurations b, c, and d are respectively of level 1 types 0000, 1111, and 1212.³²

Levels $j > 1$ use a slightly less complicated representation. As mentioned before, in practice ABACUS only needs to construct states for which $M_j \sim O(1)$ for $j > 1$, and we therefore use only two sectors for these levels. As illustrated in Fig. 3, sector 0 at level j contains allowable quantum numbers $I' \geq 0$, while sector 1 contains $I' < 0$. The *level j type* is defined similarly to the level 1 type, but is here given by the simple combination $e_1^j e_0^j$.

The whole set of level j types for $j=1, \dots, N_s$ defines the *type* of the eigenstate.

Mapping to Young tableaux

Within a given sector at a given level, we map the configuration of occupied and unoccupied quantum numbers onto a Young tableau. This is done differently in each of the four sectors of level

FIG. 3. The two sectors for levels $j > 1$. The level j type for this configuration is 13.

1, and in the two sectors of each level $j > 1$.

It is simplest to illustrate the general idea using a level $j > 1$. This is done in Fig. 4, where a type 22 configuration is considered. Sector 0 of that level has 2 excitations; there are six allowable quantum numbers. The lowest energy configuration of this type would have the two quantum numbers $1/2$ and $3/2$ occupied. This would be mapped to an empty tableau. Instead, in the configuration shown, the rightmost excitation sits on $11/2$, four slots higher than its quantum number of $3/2$ in the lowest configuration. This displacement is represented by putting four boxes in the first row of the tableau. Similarly, the second-rightmost excitation sits on $5/2$, two slots higher than its original $1/2$. We therefore put two boxes in the second row of the tableau. In sector 1, we have two excitations, and we use the same logic but inverted right to left. The first row of the tableau in sector 1 thus represents the left displacement of the excitation which sits leftmost in the lowest configuration. It should be more or less immediately clear to the reader that the generalized Pauli principle means that the tableaux that are thus obtained obey the Young tableaux rules. The tableaux in a given sector can also be read in a dual manner: if the row lengths represent, e.g., the right displacements of a set of particles/holes, then the column lengths represent the left displacements of the associated holes/particles.

At level $j=1$, we need four tableaux. The sector 0 tableau at level 1 represents the inward left displacements (toward the middle) of the holes on the right of the base Fermi set. The sector 1 tableau similarly represents the right displacements of the holes on the left of the base Fermi set. Sectors 2 and 3 represent, respectively, the right (left) displacements of the excitations to the right (left) of the base Fermi set, similarly to the logic used in levels $j > 1$.

For a given sector i at a given level j , the number of rows n_r in the corresponding tableau is equal to the number of excitations e_i^j in this sector. The number of columns n_c is given by the maximal possible displacement of the “highest” excitation, which is simply the number of available quantum numbers n_i^j in this sector minus the number of excitations. There are thus $d_i^j = \binom{n_i^j}{e_i^j} \equiv B_{n_i^j, e_i^j}$ different allowable tableaux (in other words, configurations) in this sector at this level.

We number Young tableaux in the following manner: the empty tableau is given $id=0$. We add boxes on the first row until the row is full, increasing the id by one for each added box. When we reach the maximal length, the next tableau is the one with a single box on the second row (and thus also on the first row). We then add boxes again on the first row up to the maximal length. The next tableau is then the one with two boxes on the second row. We proceed like this until both the

FIG. 4. Mapping quantum number configurations in a sector to a Young tableau, here for the two sectors at a level $j > 1$.

FIG. 5. Young tableaux and their identification numbers. Tableaux 0–9 form the set of all ten tableaux with $(n_r, n_c) = (2, 3)$. Tableaux 0–19 form the set of twenty tableaux with $(n_r, n_c) = (3, 3)$.

first and second rows are full. The next tableau is then the one with one box on the third row (and thus also on the second and first ones). This is done until the full tableau is obtained. An example of this is given in Fig. 5, where we list all the $(n_r, n_c) = (3, 3)$ tableaux and their identification numbers. It is simple to write a recursive algorithm providing the mapping to/from Young tableaux from/to identification numbers.

Eigenstate labeling in ABACUS

We now have all the needed elements to describe the labeling of eigenstates within ABACUS. An eigenstate is uniquely identified by a configuration. A configuration is uniquely defined by three integers: its base_id, its offset type_id, and its offset_id.

The base_id is a long long integer specifying the set $\{M_j\}$. In general, we have $M_1 \approx M \gg 1$, with $M_j \sim O(1)$ for $j > 1$. We thus write

$$\text{base_id} = [\#\#]_{M_{N_s}} [\#\#]_{M_{N_s-1}} \dots [\#\#]_{M_3} [\#\#]_{M_2} [\#\#]_{M_1},$$

in which $[\#\#]_{M_j}$ signifies two integers giving M_j , e.g., 03 or 11 or other two-digit non-negative number less than 100. At level 1, in order to permit the labeling of states in large systems, we allow for a five-digit specification. The numbers should be read from *right to left*, since leading zeros are scrapped by file or standard output. A type_id of 1040200293 therefore means that the base is given by $M_1=293$, $M_2=2$, $M_3=4$, $M_4=1$, and $M_j=0$, $j > 4$.

The type_id of the offset of a configuration is represented as an integer,

$$\text{type_id} = e_1^{N_s} e_0^{N_s} e_1^{N_s-1} e_0^{N_s-1} \dots e_1^3 e_0^3 e_1^2 e_0^2 e_3^1 e_2^1 e_1^1 e_0^1$$

(the numbers should again be read *from right to left*). This is also represented as a long long integer in ABACUS, since in practice the code never puts more than nine excitations in a sector (allowing the type_id to be interpreted one character at a time), and uses only the lowest set of levels (meaning that the limitations on the length of integers do not affect the actual calculations performed).

Finally, the offset_id is defined as follows. Recall that d_i^j is the total number of tableaux of sector i at level j . The tableaux identification number t_i^j for this sector at this level thus runs through the values $0, 1, \dots, d_i^j - 1$. Given a set of tableaux identification numbers $\{t_i^j\}$ for each sector at each level, we thus define

$$\text{offset_id} = \sum_{j=1}^{N_s} \sum_{i=0}^3 t_i^j \prod_{i'=0}^{i-1} d_{i'}^j \prod_{j'=1}^{j-1} \prod_{i''=0}^3 d_{i''}^{j'}$$

as the single integer defining all tableaux identification numbers. In other words, we order the (s, l) sector/level pairs as

$$(0,0),(1,0),(2,0),(3,0),(0,1),(1,1),(0,2),(1,2), \dots, (0, N_s - 1), (1, N_s - 1),$$

and the `offset_id` is obtained by multiplying each tableau id by the product of the number of possible tableaux at each lower sector/level pair.

We therefore have defined a mapping between three integers and the whole set of quantum numbers of an eigenstate,

$$(\text{base_id}, \text{type_id}, \text{offset_id}) \Leftrightarrow \{P_\alpha^j\},$$

where the string content and excitation pattern are clearly (i.e., humanly) legible from the first and second integers, respectively. These numbers are used in all functions used by ABACUS and in the data files that it produces.

IV. CONSTRUCTING INDIVIDUAL EIGENSTATES AND CALCULATING THEIR FORM FACTORS

A. Solving the Bethe equations

As described above, the Bethe ansatz provides an economical pathway for obtaining eigenfunctions by replacing the diagonalization of the Hamiltonian by the process of solving some sets of nonlinear coupled equations for one eigenstate at a time. The overall difficulty of solving the Bethe equations varies greatly from one model to the other. Also, within one model, the solution to the Bethe equations is easy to find for some of the eigenstates, while for some other eigenstates this might represent a very difficult challenge.

The simplest situation is when all solutions to all Bethe equations are to be found in terms of real, finite rapidities. This is the case, for example, for the one-dimensional Bose gas (Lieb–Liniger model) in the repulsive regime: the proof is simple and relies on the convexity of the Yang–Yang action associated with the Bethe equations (4) on the field of real numbers (Ref. 33, see also details in Ref. 3). For Heisenberg spin chains, however, the corresponding Yang–Yang action is not convex, and complex solutions exist. It is a well-known fact that there is no (and cannot be a) good numerical method for solving coupled nonlinear systems of equations, so we have to input as much preliminary knowledge as possible in the procedure. The string hypothesis conveniently provides such preliminary information.

In order to solve the BGT equations, ABACUS mixes different methods. The simplest method, used in the initial stages of the solution process, consists in performing some simple iterations. Defining the set $\{\lambda_{\alpha,(i)}^j\}$ as the set of rapidities at iteration number i , we can proceed as follows. First, the initial set is defined as the solution to the decoupled conditions $N\theta_{\text{kin}}^j(\lambda_{\alpha,(0)}^j) = 2\pi I_\alpha^j$. Subsequent sets are defined from the recursion relation $N\theta_{\text{kin}}^j(\lambda_{\alpha,(i+1)}^j) = 2\pi I_\alpha^j - \sum_{k=1}^{N_s} \sum_{\beta=1}^{M_k} \theta_{\text{scat}}^{jk}(\lambda_{\alpha,(i)}^j, \lambda_{\beta,(i)}^k)$, possibly using some damping to help stabilize the process. Such simple iterations have the advantage that they often work, in the sense that some measure of convergence is achieved. Each iterative step is also rather quick, involving order M^2 operations. However, convergence is slow, especially when the solution is approached, and it is therefore desirable to accelerate the algorithm in various ways.

One way to converge more quickly is to track the changes in the rapidities over a certain number of iterations. This gives a flow pattern to each rapidity: if this flow is sufficiently regular, it can then be extrapolated to an infinite number of iterations in order to generate a new set of rapidities. In practice, this procedure of combining iterations (say four or five) and extrapolations substantially accelerates convergence as compared to doing only simple iterations.

A more elegant approach becomes available once the iterations have come sufficiently close to the true solution: the matrix Newton method. This has the advantage of converging quadratically, but has the disadvantages that (1) it is more computationally costly per iteration step, involving order M^3 operations, and (2) it is not always stable if the starting set of rapidities is too far from the actual solution (this instability is a possibility if the Yang–Yang action is not convex). Various run-time checks need to be made to see if this procedure is stable. One advantage of the matrix Newton method is that the Jacobian matrix which is needed, given by the matrix of derivatives of

the left-hand side of (4), is precisely the Gaudin matrix which is needed to compute the norm of the eigenstate (discussed in the next subsection), which also provides a slight acceleration of the algorithm.

ABACUS considers the Bethe equations as solved provided a further iterative or Newton step gives a new set of rapidities for which the sum of square differences with the previous set is below a fixed convergence precision threshold, the latter being set to a fixed value somewhat larger than the available numerical accuracy (machine epsilon). In the case of purely real rapidities, this is the only convergence criterion applied. For states with strings, however, we have to be more careful. Namely, solutions to the BGT equations do not always represent proper solutions to the corresponding Bethe equations, since the string hypothesis might not be verified to sufficient accuracy for that specific eigenstate. In spin chains, this occurs in many circumstances, including for eigenstates with strings having a very large rapidity, or when the magnetic field is near zero (for a more extensive discussion of deviated strings, see Ref. 28 and references therein). For eigenstates with string states, ABACUS therefore also computes the first-order string deviations by substituting the solution of the BGT equations back into the original Bethe equations, yielding a new iteration for the (now full) set of rapidities, this time including their imaginary parts explicitly. If the sum of the string deviations (defined as the difference between the patterns obtained and pure, undeformed strings) is higher than a given fixed string precision threshold, the eigenstate is deemed unusable, and its contribution is excluded from the final result to prevent corruption of the data. If these deviations are small enough, the contribution from this state will, however, be kept.

B. Calculating norms and form factors

Once the set of rapidities $\{\lambda_a^j\}$ of an eigenstate is known, it is then a matter of straightforward number crunching to obtain the norm of the eigenstate. This is given by the Gaudin–Korepin formula^{34,35} in terms of the determinant of the Gaudin matrix, the latter being defined (as mentioned above) as the derivative matrix of the Bethe equations (in the case of string states, the norm can be similarly calculated from a reduced Gaudin matrix obtained from the BGT equations, see Refs. 15 and 36). The determinant itself is then calculated using LU decomposition.

For form factors of local operators in non-nested Bethe ansatz integrable models, the algebraic Bethe ansatz provides a representation in terms of a matrix determinant, similarly to the norm. In this case, however, the matrix depends on two sets of rapidities (one for each of the bra and ket states) and takes a different form depending on which operator is considered. The explicit formulas can be found in the literature. For the Bose gas, the density operator form factor is given in Refs. 7 and 8. The field operator form factor was given as a sum of determinants in Ref. 37 and simplified to a single determinant in Ref. 17. For Heisenberg $S=1/2$ chains, the S^z and $S^\pm$ operator form factors were given in Refs. 9 and 10 for the case where string states exist at most on one side of the bracket. The general expression, valid when either (or both) of the bra and ket states have strings, was given in Refs. 15 and 38. Here again, the determinant is calculated by ABACUS using LU decomposition.

V. SCANNING THE HILBERT SPACE

Assuming that eigenstates can be classified, individually labeled, and constructed as described in the previous sections, we now come to the more challenging third item on the introduction's "wish list:" the summation over intermediate states.

For a given DSF in a given model, the perfect scanning algorithm will generate eigenstates in order of numerically decreasing absolute value of form factors. This is not trivial to achieve in the models we consider, since the intuitions we can develop relating to the solutions to the Bethe equations and the rapidity dependence of the form factors remain incomplete.

All is not lost, however. Four general but simple guiding principles can be used to define an efficient algorithm to scan the Hilbert space. The first principle is that (i) for large enough system sizes, form factors obey an approximate continuity principle, in the sense that modifying an excited state in a small way (e.g., moving just one quantum number) does not dramatically change

the form factor. This principle is ultimately related to the fact that form factors are analytic functions of the rapidities of the eigenstates involved and translates to the statement that if we find some excited state having a large contribution, we can exploit this by concentrating computational resources on the pool of states in its vicinity.

The remaining three principles associate in turn with the three integers we use to label eigenstates. In general, we observe that form factors of local physical operators between the ground state and an excited state decrease in numerical value for increasing complexity of the excited state, by which we explicitly mean that form factors become smaller if (ii) more/higher strings are included in the excited state, (iii) more particle-hole excitations are created in the lowest level, and (iv) the excitations' quantum numbers are moved further and further away from the base Fermi set defining the lowest-energy state.

These principles are not strictly true in all circumstances, and the ABACUS implementation attempts to “take them with a grain of salt” and probe sufficiently widely to gather all important contributions on the one hand, but not spend too much time on irrelevant contributions on the other.

A. Climbing the Bethe tree

The best way to visualize the scanning through the Hilbert space for the excited states is by using an analogy with another of Bethe's creations, namely, a Bethe lattice.³⁹ Starting from the lowest-energy state, one can then create particle-hole pairs in the lowest level, let these particles and holes move away from the Fermi configuration, add higher strings, let those disperse, and so on.

Explicitly, the first excited state constructed is the lowest-energy state of the appropriate subspace, with the lowest `base_id` and zero `type_id` and `offset_id`. A particle-hole pair is then created, the `type_id` becoming in turn 101, 110, 1010, and 1001. For each of these four branchings, a scan is made on the position of innermost (so sector 0 or 1) excitation. If the contributions obtained to a chosen sum coming from these states is “large enough” (whose meaning we will make explicit below), the next innermost excitation is “raised” by a unit, and a new scan is made over the innermost excitation positions. The same logic is applied recursively in all sectors of all levels of the base in current use, as illustrated in Fig. 6.

If a base, through this scheme, has given a large enough contribution, the base is then “complexified” at different levels, yielding a new family of bases on which the recursive scanning is also performed in a similar way until all large enough contributions from all bases have been exhausted. Very importantly, an ABACUS object called `Scanned_Intervals` maintains a map of the Hilbert space which shows which regions have been scanned and what average results have been obtained per scanned region. This crucial run-time information will be used subsequently, as described in Sec. V B.

B. Optimizing the climb

Defining the term “large enough” used in Sec. V A leads us to the description of the highest-level procedure in ABACUS. The driving algorithm defines a real quantity called `running_threshold` which defines a “large enough” form factor as one whose value exceeds this threshold. The initial value of the threshold is chosen such that only relatively few (of order $\sim N^2$) intermediate states are constructed on a first run of the recursive Bethe tree climb. Once this first climb has terminated, a new lower value of `running_threshold` is chosen, and a new climb is initiated.

The `Scanned_Intervals` object now fulfils its role, by (1) specifying to the scanning algorithm which parts of the Hilbert space of intermediate states can potentially yield contributions which are now “large enough” according to the new `running_threshold`, and (2) which parts have already been scanned and should therefore be avoided to prevent state double counting.

When the second recursive Bethe tree climb terminates, the `running_threshold` is again reduced by an appropriate amount, and a third recursive climb is initiated. This is continued until a user-specified maximal allowed computing time is reached, before which ABACUS interprets and

FIG. 6. Climbing the Bethe tree for a given base. If an active excitation scan gives good enough results (in the sense defined in the text), a whole new generation of scans is initiated by descending the state (again, as explained in the text) in all possible ways, and starting a new active excitation scan for each such descendent. This procedure is continued recursively, and continues until good results are exhausted. Since each eigenstate is obtained via a unique such descendency path, the full Hilbert space can be faithfully probed.

saves all the data to disk. Importantly, the `Scanned_Intervals` object is also saved, allowing a calculation which has successfully terminated to be restarted at a later stage, enabling to “polish” results if required. A parallel implementation of ABACUS also exists, based on the same principles. This whole construction naturally provides an optimal use of the available computing resources.

VI. RESULTS AND THEIR INTERPRETATION

At the end of an ABACUS run, a large collection of form factor data becomes available, composed of energy-momentum-form factor triplets, either in a file with additional state-by-state information or conveniently binned into a form factor matrix (for large-scale computations producing immense numbers of entries, which could not all be individually saved to disk). The quality of this raw set of data (see an example in Fig. 7) can be measured by exploiting various sum rules, the simplest of which is to consider the summation over all momenta and integration over all frequencies of Eq. (2). The resulting local expectation value $\langle \mathcal{O}^a \mathcal{O}^{\bar{a}} \rangle$ is in most cases known analytically. Since expression (3) is a summation over terms strictly greater than or equal to zero, the form factor data set can be summed up and compared to the analytical expectation value. A sum rule saturation percentage can thus be assigned to each data set produced. Other sum rules can occasionally be used, for example, the so-called f -sum rule which relates the integral over all frequencies of $\omega S^{aa}(k, \omega)$ to a known function of k . This is particularly useful since it allows to individually evaluate the quality of each separate momentum slice in the data set, which is beyond the reach of the integrated intensity. Sum rules are routinely checked by ABACUS: a summary file is produced, which contains information about the run. For each base and associated type, the number of scanned states is given, together with their contribution to the integrated intensity sum rule. Such summary files therefore provide a wealth of information about which types of excitations carry significant correlation weight and which do not.

Inevitably, for a given target sum rule saturation, the computation quickly becomes more intensive with increasing system size. The operation count for obtaining an individual state and its form factor scales like N^3 . In practice, the number of states that has to be included in the sum-

Form factor results in a typical ABACUS run

FIG. 7. (Color online) Numerical value of (squared) form factors obtained during a typical ABACUS run (here, for the XXZ model at $\Delta=0.6$ and $S_{\text{tot}}^z=0.1$ for a small system of $N=50$ sites). The horizontal axis is simply the ordinality of the computed point. The continuous (in color: yellow) line represents a hypothetically achieved perfect ordering of form factors in monotonically decreasing numerical value (an ideal algorithm would simply follow the yellow line). While still far from perfect (a factor of approximately 10 in speed could be gained by making the scanning perfect), ABACUS does follow such ideal lines reasonably closely in most circumstances.

mation over intermediate states superficially seems to scale polynomially with N as N^{-3-5} depending on the system and its parameters. As a rule of thumb, doubling the system size thus requires about two orders of magnitude more computational power to achieve the same saturation. In fact, the scaling with system size is a complex affair. If this scaling was of a well-defined degree, the curves of leftover sum rule weight as a function of number of contributions summed would have the same slope (assuming the efficiency of ABACUS to be system size independent). This is not the case, as can be seen from the example in Fig. 8 and Table I, and there seems to be some additional (possibly logarithmic) system size dependence in the exponent. This is, in fact, quite natural: for larger and larger systems, we expect an increasing fraction of the correlation weight to be carried by states with higher excitation numbers. The system size scaling is discussed somewhat more in Sec. VII, but a systematic study of it is beyond the scope of this article.

The finite system DSF which ABACUS calculates, as is clear from Eq. (3), is not a smooth function in the energy-momentum plane, since it is simply composed of isolated delta peaks of various heights, aligned on well-defined lines in momentum but distributed in a very irregular pattern in energy (due to the adopted periodic boundary conditions on the wave functions, momentum is always an integer multiple of $2\pi/N$; energies are on the other hand certainly not equispaced in our interacting models). Only in the thermodynamic limit does the DSF become a continuous function; to obtain this, some form of smoothing of the ABACUS results is necessary. Since the momentum is already sitting on a regular lattice, only the energy delta function is smoothened into a Gaussian. This has the unfortunate effect of blurring some would-be sharp excitation thresholds, but allows to obtain easily interpretable density plots of the DSF. The width of the Gaussian is chosen so as to be somewhat larger than the typical two-excitation energy level spacing, so that the delta functions blend into a smooth function representing the density of states in this region of the Brillouin zone. Since this two-excitation energy level spacing is typically of order $1/N$, the negative effects of the Gaussian quickly become less significant for bigger systems.

FIG. 8. (Color online) Again for the typical example of the XXZ model at $\Delta=0.6$ and $S_{\text{tot}}^z=0.1$, the remainder of total integrated intensity sum rule after summing finite numbers of form factors is plotted. The number of states needed to achieve a required saturation increases more rapidly than polynomially in system size (as seen by the decreasing slope of the curves for increasing N), pointing to additional weak system size dependence in the ABACUS operation count exponent associated with the system size (see text for further discussion).

VII. DISCUSSION

A. Size dependence of contributions

The dependence on system size of a given DSF in a given model is very complex. First of all, it may take different forms depending on which model parameters are used and on which part of the energy-momentum plane one is studying. Changing system size affects (i) the number of eigenstates in the Hilbert space, (ii) the relative number of states of different bases, (iii) the quality of numerical solutions to the Bethe equations, in particular, the relative number of acceptable and unacceptable quality string states that can be constructed, (iv) the form factor of a given identifiable state, and finally (v) the positional distribution of eigenstates in energy (density of states).

Point (i) is trivial: the dimensionality of the Hilbert space is factorial in system size and poses a severe limitation on the attainable size. Point (ii) is more subtle and interesting. Namely, the

TABLE I. Number of form factors needed to attain a given sum rule saturation level for various values of system size. This is again for the example of the XXZ model at $\Delta=0.6$ and $S_{\text{tot}}^z=0.1$. The numbers in parentheses are the Hilbert space dimensionality fraction represented by the number of states mentioned, illustrating the efficiency of working in the chosen eigenstates basis. Missing numbers in the table could be obtained by simply running ABACUS for longer (all data here are from single CPU runs). Although the number of states needed to achieve a fixed target sum rule saturation increases more rapidly than polynomially in N as seen from Fig. 7, the Hilbert space fraction needed also decreases.

%	$N=50$	$N=100$	$N=200$	$N=400$
90	235(4.9×10^{-12})	944(6.8×10^{-26})	8 772(5.3×10^{-54})	117 373(3.5×10^{-111})
95	385(8.2×10^{-12})	2 492(1.8×10^{-25})	27 096(1.6×10^{-53})	573 466(1.7×10^{-110})
99	1 515(3.2×10^{-11})	18 490(1.3×10^{-24})	469 120(2.8×10^{-52})	
99.9	7 478(1.6×10^{-10})	141 560(1.0×10^{-23})		
99.99	26 724(5.7×10^{-10})	932 706(6.8×10^{-23})		

number of n -excitation states scales approximately as (system size) n ; therefore, the number of $2, 3, 4, \dots$ excitation states scales differently, families with more excitations having an increasingly larger number of members. For any n -excitation state, we can generically find order (size) 2 states with $n+2$ excitations, and determining how these different bases relatively contribute to the final answer for large systems is thus an extremely complicated problem.

We can, however, make some general statements. Suppose for definiteness that our excitations only come in pairs (as in, e.g., the Lieb–Liniger model). For very small size (few particles), two-excitation states will carry a fraction $c_2 \sim 1$ of the sum rules. As the system size N increases, we will see the two-excitation contribution “leak” (as guaranteed by the preservation of sum rules) into the increasingly more numerous $4, 6, 8, \dots$ -excitation ones. c_2 will be strictly monotonically decreasing in N , e.g., $\Delta c_2 / \Delta N < 0 \forall N$. Despite this, the limit $\lim_{N \rightarrow \infty} c_2$ might remain finite (as is the case for spinons in the zero-field isotropic or gapped Heisenberg antiferromagnet; for the gapless region, it seems reasonable for c_j for j finite to vanish in the thermodynamic limit, except at zero field). c_4 , on the other hand, will not be monotonic: it will increase for small N but eventually reach a maximum at some inflexion value N_4 , after which it will decrease, i.e., $\Delta c_4 / \Delta N > 0$, $N < N_4$ but $\Delta c_4 / \Delta N < 0$, $N > N_4$. The limiting value $\lim_{N \rightarrow \infty} c_4$ might also remain finite. There will similarly be increasingly large inflexion values $N_6, N_8, \dots$, $N_4 \ll N_6 \ll N_8 \dots$, but the determination of these inflexion values is beyond the reach of the current implementation.

This scenario, although more or less inevitable in the mind of the author, remains conjectural since even the lowest of these inflexion points (N_4) seems to lie above currently achievable system sizes. In any case, we can expect the finite-size result for large enough systems to closely resemble the infinite-size result (except perhaps in the immediate vicinity of excitation thresholds): while each term in the series $c_2 + c_4 + c_6 + \dots$ might have substantial size dependence, the sum itself (i.e., the full DSF) is seen to have much smaller size dependence: increasing contributions (as system size increases) from higher excitation numbers tend more or less to compensate the decrease in the lower excitation number contributions (at least in gapless models), since these two types of contribution resemble each other. Thus, for example, the contribution of four-spinon states in the Heisenberg magnet⁴⁰ looks like a rescaled two-spinon contribution, and we can expect more or less the same of yet higher excitation numbers (this, disregarding the different high-energy tails, which carry very little correlation weight anyway).

VIII. CONCLUSION AND PERSPECTIVES

Much remains to be done in the general field of dynamical correlation functions of integrable models, and the reader is referred to Ref. 20 for an extensive discussion of this subject. As far as ABACUS is concerned, various improvements of the subalgorithms for solving BGT equations are possible, as well as further optimization of the state scanning algorithm. A better handling of deformed string eigenstates is also necessary: this, at the moment, represents the most severe limitation of the applicability of ABACUS for some DSFs of spin chains in small magnetic fields, but requires a rather more elaborate treatment to be properly dealt with. For certain cases, as, for example, gapped antiferromagnets, the correct state counting itself is not fully known for arbitrary bases due to difficulties in defining the quantum number limits (see Ref. 41 for further discussion of this point).

The restriction to zero temperature would also need to be overcome, possibly by using some ideas from the thermodynamic Bethe ansatz. On a more theoretical front, representations of form factors for nested systems such as the Yang permutation model⁴² or the Hubbard model⁴³ are needed before ABACUS can be used. Once these are available, however, the “universal” ABACUS logic will be quickly implementable for these systems as well.

A. Data availability

Requests to the author for calculations of DSFs for specific models and values of the parameters are welcome. An online database of ABACUS computations will also eventually be made available. Alternately, the C++ source code can be obtained from the author on request.

ACKNOWLEDGMENTS

Much work went into the development of ABACUS, which was initiated during a collaboration with J. M. Maillet. Subsequent versions of the spin chain part of the code were checked against an independent implementation by my former student R. Hagemans. In a similar fashion, the one-d Bose gas part of the code was proof checked against an independent implementation by P. Calabrese.

- ¹H. Bethe, *Z. Phys.* **71**, 205 (1931).
- ²M. Gaudin, *La Fonction d'Onde de Bethe* (Masson, Paris, 1983), and references therein.
- ³V. E. Korepin, N. M. Bogoliubov, and A. G. Izergin, *Quantum Inverse Scattering Method and Correlation Functions* (Cambridge University Press, Cambridge, 1993), and references therein.
- ⁴M. Takahashi, *Thermodynamics of One-Dimensional Solvable Models* (Cambridge University Press, Cambridge, 1999), and references therein.
- ⁵B. Sutherland, *Beautiful Models* (World Scientific, Singapore, 2004), and references therein.
- ⁶T. Giamarchi, *Quantum Physics in One Dimension* (Oxford University Press, Oxford, 2004), and references therein.
- ⁷N. A. Slavnov, *Theor. Math. Phys.* **79**, 502 (1989).
- ⁸N. A. Slavnov, *Theor. Math. Phys.* **82**, 273 (1990).
- ⁹N. Kitanine, J. M. Maillet, and V. Terras, *Nucl. Phys. B* **554**, 647 (1999).
- ¹⁰N. Kitanine, J. M. Maillet, and V. Terras, *Nucl. Phys. B* **567**, 554 (2000).
- ¹¹D. Biegel, M. Karbach, and G. Müller, *Europhys. Lett.* **59**, 882 (2002).
- ¹²D. Biegel, M. Karbach, and G. Müller, *J. Phys. A* **36**, 5361 (2003).
- ¹³J. Sato, M. Shiroishi, and M. Takahashi, *J. Phys. Soc. Jpn.* **73**, 3008 (2004).
- ¹⁴J.-S. Caux and J. M. Maillet, *Phys. Rev. Lett.* **95**, 077201 (2005).
- ¹⁵J.-S. Caux, R. Hagemans, and J. M. Maillet, *J. Stat. Mech.*, 09003 (2005).
- ¹⁶J.-S. Caux and P. Calabrese, *Phys. Rev. A* **74**, 031605(R) (2006).
- ¹⁷J.-S. Caux, P. Calabrese, and N. Slavnov, *J. Stat. Mech.: Theory Exp.*, 01008 (2007).
- ¹⁸J.-S. Caux (unpublished).
- ¹⁹J.-S. Caux (unpublished).
- ²⁰J.-S. Caux, J. M. Maillet and A. Tennant (unpublished).
- ²¹For example, for Heisenberg spin chains, where the $\hat{z}$ component of the total magnetization S_{tot}^z is conserved, M would be the number of overturned spins starting from a fully ferromagnetic reference state.
- ²²This is correct as long as a degree of flexibility is maintained in what we mean with a wavefunction of Bethe ansatz form. An example that comes to mind is the Heisenberg spin chain at roots of unity (Refs. 44 and 45), where a complete set of states includes wavefunctions which can be viewed as given by the *derivative* of a Bethe ansatz.
- ²³M. Gaudin, *Phys. Rev. Lett.* **26**, 1301 (1971).
- ²⁴M. Takahashi, *Prog. Theor. Phys.* **46**, 401 (1971).
- ²⁵M. Takahashi, *Prog. Theor. Phys.* **48**, 2187 (1972).
- ²⁶In Bethe ansatz as in many other things, the devil is in the details: the generalized Pauli principle is not strictly valid (see cited literature in Ref. 28) in its most basic formulation, but can be made so by an appropriate cataloguing of string states, including (over)deformed strings.
- ²⁷While in most cases the left- and right-limiting quantum numbers coincide, this is not always true. The simplest example of such a case is the isotropic Heisenberg chain, for which one should take $I_{\infty,+}^1 = I_{\infty,-}^1 + 1$ or the converse, since the associated rapidity then sits on ∞ , which is not distinguished from $-\infty$. Such extra states also occur, for example, in the gapped $S=1/2$ antiferromagnet and in the gapless XXZ chain at roots of unity. The states obtained from these exceptional solutions to the Bethe equations are associated with certain discrete symmetries of the model [the global $SU(2)$ symmetry in the case of the isotropic chain, and the loop algebra symmetry for the gapless chain at roots of unity] and can be treated in the ABACUS logic by, for example, introducing additional levels.
- ²⁸R. Hagemans and J.-S. Caux, *J. Phys. A: Math. Theor.* **40**, 14605 (2007).
- ²⁹Such pictorial representations of quantum numbers and of their movement when scanning through the set of eigenstates are reminiscent of a traditional ABACUS and suggested the name of the library.
- ³⁰L. D. Faddeev and L. A. Takhtajan, *Phys. Lett. A* **85**, 375 (1981).
- ³¹We assume that ABACUS will not need to construct states with extremely many excitations, and we therefore restrict the e_i^l to $0 \leq e_i^l \leq 9$, making the above labeling unambiguous.
- ³²For the specific case of the one-dimensional Bose gas, a different convention is used, in which sectors $1/3$ and $2/0$ are permuted.
- ³³C. N. Yang and C. P. Yang, *J. Math. Phys.* **10**, 1115 (1969).
- ³⁴M. Gaudin, *J. Math. Phys.* **12**, 1674 (1971); **12**, 1677 (1971).
- ³⁵V. E. Korepin, *Commun. Math. Phys.* **86**, 391 (1982).
- ³⁶A. N. Kirillov and V. E. Korepin, *J. Math. Sci.* **40**, 13 (1988).
- ³⁷T. Kojima, V. E. Korepin, and N. A. Slavnov, *Commun. Math. Phys.* **188**, 657 (1997).
- ³⁸There is an unfortunate typo in Eq. (29), where L should be defined with a minus sign.
- ³⁹H. A. Bethe, *Proc. R. Soc. London, Ser. A* **150**, 552 (1935).
- ⁴⁰J.-S. Caux and R. Hagemans, *J. Stat. Mech.: Theory Exp.*, 12013 (2006).
- ⁴¹J.-S. Caux, J. Mossel, and I. Pérez Castillo, *J. Stat. Mech.: Theory Exp.*, 08006 (2008).
- ⁴²C. N. Yang, *Phys. Rev. Lett.* **19**, 1312 (1967).

⁴³F. H. L. Essler, H. Frahm, F. Göhmann, A. Klümper, and V. E. Korepin, *The One-Dimensional Hubbard Model* (Cambridge University Press, Cambridge, 2005), and references therein.

⁴⁴K. Fabricius and B. McCoy, *J. Stat. Phys.* **103**, 647 (2001).

⁴⁵K. Fabricius and B. McCoy, *J. Stat. Phys.* **104**, 573 (2001).