



Universiteit
Leiden

The Netherlands

Fermions, criticality and superconductivity

She, J.H.

Citation

She, J. H. (2011, May 3). *Fermions, criticality and superconductivity*. *Casimir PhD Series*. Faculty of Science, Leiden University. Retrieved from <https://hdl.handle.net/1887/17607>

Version: Corrected Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/17607>

Note: To cite this publication please use the final published version (if applicable).

CHAPTER 3

FERMIONS IN THE CONSTRAINED PATH INTEGRAL: TOWARDS THE MINUS SIGN PROBLEM

3.1 Introduction

The ‘quantum weirdness’ of the Fermi-gas is obvious: how to understand the Fermi-surface, the Fermi-energy and so forth, just knowing about classical statistical physics? The interacting Fermi-liquid is a bit more than the Fermi-gas, but focusing on the emergence principles it is deep inside the same thing. As Landau pointed out, the Fermi-liquid is connected by adiabatic continuation to the Fermi-gas meaning that the two are qualitatively indistinguishable at the long times and distances where emergence is in full effect. The great framework of diagrammatic perturbation theory developed in the 1950’s [27] does allow to arrive at quite non trivial statements associated with the presence of the interactions but it only works under the condition that the Fermi-liquid is adiabatically connected to the Fermi gas. But conventional Feynman diagrams are impotent with regard to revealing the nature of ‘non Fermi liquids’. To complete the ‘fermionic’ repertoire of theoretical physics, Bardeen, Cooper and Schrieffer discovered the ‘Hartree-Fock’ mechanism, showing how the Fermi-gas can become unstable towards a bosonic state, like the superfluids- and conductors, charge- and spin density wave states and so forth. Despite fermionic peculiarities (like the gap function), this is eventually a recipe telling us how the fermi-gas can

turn into bosonic matter that is in turn ruled by the Ginzburg-Landau-Wilson classical emergence rules.

Given the present repertoire of theoretical physics, all we know to do with fermionic matter is to hope that it is a Fermi gas or bound in bosons. But we are facing a zoo of ‘non-Fermi-liquid’ states of electrons coming out of the experimental laboratories and the theorists are standing empty handed because the fermion signs render all the fancy theoretical technologies to be useless. The NP hardness of the sign problem tells us that there is no mathematically exact solution but how many features of the physical world we understand well are actually based on exact mathematics? Nearly all of it is based on an effective description, mathematics that is tractable while it does describe accurately what nature is doing although it is not derived with exact mathematics from the first principles. Is there a way to handle non-Fermi-liquid matter on this phenomenological level?

The remainder of this chapter is dedicated to the case that there is reason to be optimistic. This optimism is based on a brilliant discovery some fifteen years ago of an alternative path-integral description of the fermion problem by David Ceperley [73,74]. This ‘constrained’ or ‘Ceperley’ path integral has a Boltzmannian structure (i.e., only positive probabilities) but the signs are traded in for another unfamiliar structure: a structure of constraints acting on a ‘bosonic’ configuration space that is coding for all the effects of Fermi-Dirac statistics. This is called the reach and it amounts to the requirement that for all imaginary times τ between zero and $\hbar\beta$ ($\beta = 1/(k_B T)$) the worldline configurations should not cross the hypersurface determined by the zero’s of the full N -particle, imaginary time density matrix. Although the constrained path integral suffers from a self-consistency problem since the exact constrain structure is not known except for the non-interacting Fermi-gas, it appears that this path integral is quite powerful for the construction of phenomenological effective theories. The information carried by the reach lives ‘inside’ the functional integral and should therefore be averaged. This implies that only global- and averaged properties of this reach should matter for the physics in the scaling limit. The reach is in essence a high dimensional geometrical object, closely related to the more familiar ‘nodal hypersurface’ associated with the sign changes of ground state wave function. The theoretical program is to classify the geometrical and topological properties of the reach in general terms, to find out how this information is averaged over in the path integral, with the potential to yield eventually a systematic classification of phenomenological theories of fermionic matter.

Given that Ceperley derived his path integral already quite some time ago, why is it not famous affair? These path integral are not so easy to handle. Although various interesting results were obtained [75], even the attempt to reconstruct the Fermi-liquid in this language stalled. But these efforts were limited to a very small community, with a focus on large scale numerical calculations. The potential of the Ceperley path integral to address matters of principle appears to be overlooked in the past. We discovered the Ceperley path integral in an attempt to understand the scale invariant fermionic quantum critical states as

found in the heavy fermion intermetallics. We started out on the more primitive level of wave function nodal structure, discovering by accident the much more powerful Ceperley path integral approach. We believe that we have delivered proof of principle that this language gives penetrating insights in the nature of a prominent non-Fermi liquid state: the fermionic quantum critical states realized in the heavy fermion intermetallics. Since this work is still under review we will not address it in any detail. However, to make further progress, we were confronted with the need to better understand the detailed workings of the Ceperley path integral and we decided to revisit the description of the Fermi gas and the Fermi liquid. The outcomes of this pursuit are summarized in this chapter.

This remainder of this chapter is organized as follows. In section 3.2 we introduce the Ceperley path integral, reviewing its derivation as well as various other technical issues. Section 3.3 is intended to be the highlight of this chapter. We present a quite simple solution of the Ceperley path integral for the Fermi-gas: the Fermi-gas turns out to be in one-to-one correspondence with a system of cold atoms in an harmonic trap, subjected to a deep optical lattice potential such that the atoms form a perfect Bose Mott-insulator! Finally in section 3.4 we turn to the real space description of the Fermi-gas. The presence of the reach changes radically the winding statistics as compared to the boson case and it appears that the windings of the Ceperley particles in *any* higher dimension are counted as if they are the windings associated with soft core bosons living in one space dimension.

3.2 Ceperley's constrained path integral

In this section we review Ceperley's 1991 discovery of a path integral representation for arbitrary fermion problems that is not suffering from the 'negative probabilities' of the standard formulation [73]. Surely, one cannot negotiate with the NP-hardness of the fermion problem and Ceperley's path integral is not solving this problem in a mathematical sense. However, the negative signs are transformed away at the expense of a structure of constraints limiting the Boltzmannian sum over world-line configurations. These constraints in turn can be related to a geometrical manifold embedded in configuration space: the 'reach', which is a generalization of the nodal hypersurface characterizing wave functions to the fermion density matrix. This reach should be computed self-consistently: it is governed by the constrained path integral that needs itself the reach to be computed. This is again a NP-hard problem and Ceperley's path integral is therefore not solving the sign problem. However, the reach contains all the data associated with the differences between bosonic and fermionic matter, and only its average and global properties should matter for the physics in the scaling limit since it acts on worldline configurations that themselves are averaged. Henceforth, it should be possible in principle to classify all forms of fermionic matter in a phenomenological way by classifying the average geometrical- and

topological properties of the reach, to subsequently use this data as an input to solve the resulting bosonic path integral problem. This procedure is supposedly a unique extension of the Ginzburg-Landau-Wilson paradigm for bosonic matter to fermionic matter. We do not have a mathematical proof that this procedure will yield a complete classification of fermionic matter, but we have some very strong circumferential evidences in the offering that it will work. The status of our claim is conjectural in the mathematical sense.

Let us start out presenting the answer. Ceperley proved in 1991 that the following path integral is strictly equivalent to the standard fermion path integral Eq. (1.11,1.12), as we reviewed in the introduction,

$$\rho_F(\mathbf{R}, \mathbf{R}; \beta) = \frac{1}{N!} \sum_{\mathcal{P}, \text{even}} \int_{\gamma: \mathbf{R} \rightarrow \mathcal{P}\mathbf{R}}^{\gamma \in \Gamma_\beta(\mathbf{R})} \mathcal{D}\mathbf{R} e^{-\mathcal{S}[\mathbf{R}]/\hbar}. \quad (3.1)$$

This is quite like the standard path integral, except that one should only sum over *even* permutations (the reason to address this in section IV), while the allowed worldline configurations γ are constrained to lie ‘within the reach Γ ’. This reach is defined as,

$$\Gamma_\beta(\mathbf{R}) = \{\gamma : \mathbf{R} \rightarrow \mathbf{R}' | \rho_F(\mathbf{R}, \mathbf{R}(\tau); \tau) \neq 0\} \quad (3.2)$$

for all imaginary times $0 < \tau < \hbar\beta$. In words, only those worldline configurations should be taken into account in Eq. (3.1) that do not cause a sign change of the full density matrix at every intermediate imaginary time between 0 and $\hbar\beta$. In outline, the proof of this result is as follows. The fermion density matrix is defined as a solution to the Bloch equation

$$\frac{d\rho_F(\mathbf{R}_0, \mathbf{R}; \beta)}{d\beta} = -H\rho_F(\mathbf{R}_0, \mathbf{R}; \beta) \quad (3.3)$$

with initial conditions

$$\rho_F(\mathbf{R}_0, \mathbf{R}; \beta = 0) = \frac{1}{N!} \sum_{\mathcal{P}} (-1)^p \delta(\mathbf{R}_0 - \mathcal{P}\mathbf{R}). \quad (3.4)$$

In the following we fix the reference point \mathbf{R}_0 and define the reach $\Gamma(\mathbf{R}_0, \tau)$ as before as the set of points $\{\mathbf{R}_{\tau'}\}$ for which there exists a continuous space-time path with $\rho_F(\mathbf{R}_0, \mathbf{R}_{\tau'}; \tau') > 0$ for $0 \leq \tau' < \tau$. Suppose that the reach is known in advance. It is a simple matter to show that the problematical initial condition, Eq. (3.4), imposing the anti-symmetry can be replaced by a zero boundary condition on the surface of the reach. It follows because the fermion density matrix is a unique solution to the Bloch equation (3.3) with the zero boundary condition. One can now find a path integral solution without the minus signs. One simply restricts the paths to lie in the reach $\Gamma(\mathbf{R}_0, \tau)$ imposing the zero boundary condition on the surface of the reach. The odd permutations fall for sure out of the reach since $\rho_F(\mathbf{R}_0, \mathcal{P}_{\text{odd}}\mathbf{R}_0) = -\rho_F(\mathbf{R}_0, \mathbf{R}_0)$.

The Ceperley path integral revolves around the reach. How to think about this object? The way the path integral is constructed seems to break imaginary time translations. One has to first pick some ‘reference point’ \mathbf{R} in configuration space at imaginary time 0 or $\hbar\beta$. Starting from this set of particle coordinates, one has to spread them out in the form of worldline configurations to check at every time slice that the density matrix does not change sign. The dimensionality of the density matrix is $2dN + 1$ (twice configuration space plus a time axis) and the dimensionality of the reach is therefore $2dN$ (one overall constraint). However, when we first pick a reference point \mathbf{R} and we focus on a particular imaginary time the dimensionality of this restricted reach is $dN - 1$. In the limit $\tau \rightarrow \infty$ this restricted reach turns into a more familiar object: the nodal hypersurface associated with the ground state wave function. The density matrix becomes for a given \mathbf{R} in this limit,

$$\rho(\mathbf{R}, \mathbf{R}'; \beta = \infty) = \Psi^*(\mathbf{R})\Psi(\mathbf{R}') \quad (3.5)$$

and the zero's of the density matrix are just coincident with the nodes of the ground-state wave function, $\Psi(\mathbf{R}) = 0$, where we have assumed that the ground state is non-degenerate. The wave function is anti-symmetric in terms of the fermion coordinates,

$$\Psi(\cdots, \mathbf{r}_i, \cdots, \mathbf{r}_j, \cdots) = -\Psi(\cdots, \mathbf{r}_j, \cdots, \mathbf{r}_i, \cdots), \quad (3.6)$$

and therefore the nodal hypersurface

$$\Omega = \{\mathbf{R} \in \mathbb{R}^{Nd} | \Psi(\mathbf{R}) = 0\} \quad (3.7)$$

is a manifold of dimensionality $\dim\Omega = Nd - 1$ embedded in Nd -dimensional configuration space. This nodal surface Ω is surely an object that is simpler than the full reach Γ and it is rather natural to train the intuition using the former. According to Ceperley's numerical results [73], it appears that at least for the Fermi gas the main features of the reach are already encoded in Ω . In a way, the dependence on imaginary time is remarkably smooth and unspectacular. A greater concern is the role of the reference point, or either the fact that the reach depends on two configuration space coordinates. In the long imaginary time limit, the reach factorizes in the nodal surfaces (Eq. (3.5)), which means that one can get away just considering the nodal surface of the ground state wave function, but this is not the case at finite imaginary times. It is not at all that clear what role the ‘relative distance’ $\mathbf{R} - \mathbf{R}'$ plays, although there is some evidence that it can be quite important as we will discuss in Section IX. Notice that the conventional ‘fixed-node’ quantum Monte-Carlo methods aim at a description of the ground state, using typically diffusion Monte-Carlo methods. As input for the ‘fermionic-side’, these only require the wave function nodal structure. The difference between the reach and this nodal structure is telling us eventually about the special nature of the excitations in the fermion systems since

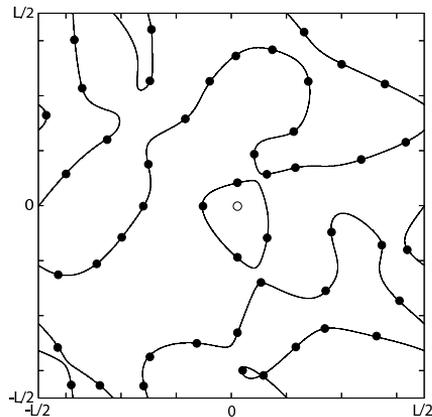


Figure 3.1: Cut through the nodal hypersurface of the ground-state wave function of $N = 49$ free, spinless fermions in a two-dimensional box with periodic boundary conditions. The cut is obtained by fixing $N - 1$ fermions at random positions (black dots) and moving the remaining particle (white dot) over the system. The lines indicate the zeros of the wave function (nodes). Note that the nodal surface cut has to connect the $N - 1$ fixed particles since the Pauli surface is a lower dimensional submanifold of dimension $Nd - d$ included in the nodal hypersurface with dimension $Nd - 1$.

the Ceperley path integral can be used to calculate dynamics, either in the form of finite temperature thermodynamics or, by Wick rotation to real time, about dynamical linear response. At this moment in time it is not well understood what the precise meaning is of these ‘dynamical signs’ encoded in the non-local nature of the reach.

Another useful geometrical object associated with Fermi-Dirac statistics is the Pauli surface, corresponding with the hypersurface in configuration space where the wave function vanishes because the fermions are coincident in real space,

$$\begin{aligned}
 P &= \bigcup_{i \neq j} P_{ij} \\
 P_{ij} &= \{\mathbf{R} \in \mathbb{R}^{Nd} \mid \mathbf{r}_i = \mathbf{r}_j\}.
 \end{aligned}
 \tag{3.8}$$

Obviously, the Pauli surface is a submanifold of the nodal hypersurface of dimension $\dim P = Nd - d$. The specialty of one dimension is that the Pauli- and nodal hypersurfaces are coincident. This property that the nodes are ‘attached’ to the particles is the key to the special status of one dimensional physics as we will explain in detail in the next section.

In the next sections we will discuss in more detail the few facts that are known about the reach and nodal hypersurface geometry and topology. To complete the

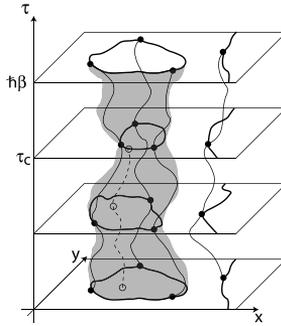


Figure 3.2: Nodal constraint structure in space-time seen by one particular particle. In the constraint path integral only world-line configurations $\{\mathbf{r}_\tau\}$ are allowed that do not cross or touch a node of the density matrix on all time slices, $\rho_F(\mathbf{R}_0, \mathbf{R}_\tau, \tau) \neq 0$ for $0 \leq \tau < \hbar\beta$. Therefore, a particular particle (white circle) is constrained by the dynamical nodal tent (grey surface) spanned by the $N - 1$ remaining particles trajectories (black circles). In a Fermi liquid the nodal tent has a characteristic dimensions and particles feel the nodal constraints at an average time scale τ_c . Later we will see that these scales are in one-to-one correspondence with the Fermi degeneracy scale E_F .

discussion of the basic structure of the Ceperley Path Integral, let us once more emphasize that according to its definition Eq. (3.1) one still has to sum over *even* permutations in so far these do not violate the reach. As for the signful path integral, this translates via the sum over cycles into a sum over winding numbers that are now associated with triple exchanges of particles. We explained already in detail in section IV that this has the peculiar consequence that it codes for supersymmetry when one is dealing with the free quantum gas that just knows about the even permutation requirement. Because of the constraints, the ‘particles’ of the Ceperley path integral are actually very strongly interacting and it is unclear to what extent this supersymmetry is of any relevance to the final solution. In fact, we do know for the Fermi-gas that the combined effect of the constraints and the triple exchanges is to eventually give back a free gas with Fermi-Dirac statistics. As we discussed in section IV, there is a ‘don’t worry theorem’ at work because the thermodynamics of the supersymmetric gas is quite similar to the Bose gas.

In conclusion, Ceperley has demonstrated that in principle fermion problems can be formulated in a probabilistic, Boltzmannian mathematical language, paying the prize of a far from trivial constraint structure that is a-priori not known while it cannot be exactly computed. Qualitatively, the reach is like the nodal structure of a wave function. It is obvious that the nodal structure codes for physics but this connection is largely unexplored, while the remainder of this

chapter is dedicated to the case that it is actually quite easy to make progress, at least with regard to the Fermi-liquid. One particular property is so important that it should be already introduced here. Any wave function of a system of fermions has the anti-symmetry property Eq. (3.6) and naively one could interpret this as ‘any physical system of fermions has its fermionic physics encoded in a $Nd - 1$ dimensional nodal surface’. This is obviously not the case. It is easy to identify a variety of fermionic systems where many more nodes are present in the fermion wave function than are required to encode the physics. A first example are Mott-insulating antiferromagnets on bipartite lattices. Because the electrons are localized they become effectively distinguishable. One can therefore transform away remnant signs in the Heisenberg spin problem by Marshall sign transformations: the bottom line is that such Mott-insulators can be handled by standard bosonic quantum Monte Carlo methods. A next example is physics in one dimensions, as we will discuss in the next section, where again the fermion signs can be transformed away completely, in a way that can be neatly understood in terms of the topology of the nodal surface. Nodal structure is therefore like a gauge field: it carries redundant information that is inconsequential for the physics. Nodal structure that is in this ‘gauge volume’ we call *reducible* nodal structure, while the ‘gauge invariant’ (physical) part of the nodal structure we call *irreducible*, and as a first step one should always first isolate the true, irreducible signs.

3.3 The Fermi gas as a cold atom Mott-insulator in momentum space

The Fermi-gas of the canonical formalism is very easy to solve exactly, and one would expect that in one or the other way this should mean that the constrained path integral is also easy to solve. This is not true at all in the position representation, as we will discuss in the next section. However, considering the derivation of the Ceperley path integral there is actually no preferred status of real space. The construction is completely independent of the representation one chooses for the single particle states. On the canonical side momentum space is the convenient representation to start from in the galilean continuum, or either any other basis that diagonalizes the single particle problem. As we will show in this section, also the Ceperley path integral of the Fermi-gas becomes very easy indeed when one chooses to formulate it in momentum space. After a couple of straightforward manipulations one finds a sign free, Boltzmannian path integral showing a most entertaining correspondence: the Fermi-gas is in one-to-one correspondence with a system of classical atoms forming a Mott insulating state in the presence of a commensurate optical lattice of infinite strength, living in a harmonic potential trap of finite strength (see Fig. 3.3a). This is literal and the only oddity is that this trap lives in momentum space instead of real space; the Fermi surface is just the boundary between the occupied optical lattice sites and

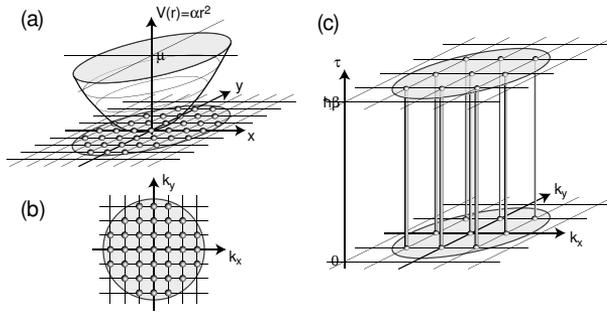


Figure 3.3: (a) a system of classical atoms forming a Mott insulating state in the presence of a commensurate optical lattice of infinite strength, living in a harmonic potential trap $V(\vec{r}) = \alpha r^2$ of finite strength; (b) the trap in momentum space k_x, k_y instead of real space; the Fermi surface is just the boundary between the occupied optical lattice sites and the empty ones; (c) a grid of allowed momentum states $k = (2\pi/L)(k_x, k_y, k_z, \dots)$ where the k_i 's are the usual integers and any worldline just closes on itself along the imaginary time τ direction $0 \rightarrow \beta$: single particle momentum conservations prohibit anything but the one cycles.

the empty ones. This boundary is sharp at zero temperature but it smears at finite temperature because of the entropy that can be gained by exciting atoms out of the trap! When you are quick, you should already have realized that this trap interpretation is actually consistent with everything we know about the Fermi-gas. Let us now proof it by constructing the Ceperley path integral.

The central wheel of the Ceperley path integral is the fermion density matrix. One should first guess an ansatz, use it to construct the path integral, to check if the same density matrix is produced by the path integral. Surely we know the full fermion density matrix for the Fermi gas, and in momentum space this turns out to be a remarkably simple affair. The k -space density matrix can be written as the determinant formed from imaginary time single particle propagators in the galilean continuum,

$$g(\mathbf{k}, \mathbf{k}'; \tau) = 2\pi\delta(\mathbf{k} - \mathbf{k}')e^{-\frac{|\mathbf{k}|^2\tau}{2\hbar M}}. \quad (3.9)$$

Since we live in the space of exact single quantum numbers these propagators are diagonal; in the galilean continuum this just means the conservation of momentum, but when translational symmetry is broken one should use here just the basis diagonalizing the single particle Hamiltonian.

Consider now the full momentum configuration space $\mathbf{K} = (\mathbf{k}_1, \dots, \mathbf{k}_N)$ imaginary time density matrix,

$$\rho_F(\mathbf{K}, \mathbf{K}'; \tau) = \frac{1}{N!} \sum_{\mathcal{P}} (-1)^{\mathcal{P}} \prod_{i=1}^N g(\mathbf{k}_{\mathcal{P}(i)}, \mathbf{k}'_i; \tau). \quad (3.10)$$

We find that the delta functions cause a great simplification. Substituting the single-fermion expression Eq. (3.9) in this expression for the density matrix Eq. (3.10) we obtain:

$$\begin{aligned} \rho_F(\mathbf{K}, \mathbf{K}'; \tau) &= \frac{1}{N!} e^{-\sum_{i=1}^N \frac{|\mathbf{k}_i|^2 \tau}{2\hbar M}} \\ &\times \sum_{\mathcal{P}} (-1)^{\mathcal{P}} \prod_{i=1}^N 2\pi \delta(\mathbf{k}_{\mathcal{P}(i)} - \mathbf{k}'_i). \end{aligned} \quad (3.11)$$

Since the single particle propagators are eigenstates of the Hamiltonian, the momentum world lines go ‘straight up’ in the time direction until they arrive at the time τ where the reconnections can take place associated with the permutations. But the δ function enforces that the permuted momentum has to be the same as the non-permuted one, and the worldlines can therefore not wind except when the momenta of some pairs of fermions coincide. But now the sum of the permutations in Eq. (3.11) is zero due to the Pauli principle. Mathematically, this follows from the fact that the expression on the right hand side of Eq. (3.11) is actually a Slater determinant formed from the delta-functions $2\pi\delta(\mathbf{k}_{\mathcal{P}(i)} - \mathbf{k}'_i)$ as the matrix elements of the $Nd \times Nd$ matrix, that are indexed by momenta $\{\mathbf{k}_{\mathcal{P}(i)}, \mathbf{k}'_i\}$. Hence, when two of the momenta coincide (e.g. $\mathbf{k}_i = \mathbf{k}_j$, $i \neq j$) there are two coinciding rows/columns in the matrix and the Slater determinant equals zero. The result is that Eq. (3.10) factorizes in $N!$ relabeling copies, associated with $N!$ nodal cells like in 1+1D, of the following simple density matrix describing distinguishable and localized particles in momentum space,

$$\rho_F(\mathbf{K}, \mathbf{K}'; \tau) = \prod_{\mathbf{k}_1 \neq \mathbf{k}_2 \neq \dots \neq \mathbf{k}_N}^N 2\pi \delta(\mathbf{k}_i - \mathbf{k}'_i) e^{-\frac{|\mathbf{k}_i|^2 \tau}{2\hbar M}}. \quad (3.12)$$

This has the structure of a Boltzmannian partition sum of a system subjected to steric constraints: it is actually the solution of the Ceperley path integral for the Fermi gas in momentum space! Let us apply periodic boundary conditions so that on every time slice of the Ceperley path integral we find a grid of allowed momentum states $\mathbf{k}_i = (2\pi/L)(k_{i,x}, k_{i,y}, k_{i,z}, \dots)$ where the $k_{i,\alpha}$ ’s are the usual integers (see Fig. 3.3b). We learn directly from Eq. (3.12) that we can ascribe a distinguishable particle with every momentum cell, with a worldline that just closes on itself along the time direction: single particle momentum conservation prohibits anything but the one cycles (see Fig. 3.3c). In addition, we find that the reach just collapses to the Pauli hypersurface, just as in one dimensions: per momentum space cell either zero or one worldline can be present. These worldlines are given by Eq. (3.9): since we are living in exact quantum number space these just go straight up along the time direction, since there are no quantum fluctuations: these are actually classical particles living in momentum space. We do have to remember that these world ‘rods’ carry a fugacity set by a potential

$\frac{|k|^2 \tau}{\hbar M}$. Henceforth, we have a problem of an ensemble of classical hard core particles that live on a lattice of ‘cells’ in momentum space where every cell can either contain one or no particle, with an overall harmonic potential envelope centered at $\mathbf{k} = 0$: this is literally the problem of cold atoms living in a harmonic trap, subjected to an infinite strong optical lattice potential, tuned such that they form a Mott-insulating state. The ground state is simple: occupy the cells starting at $\mathbf{k} = 0$, while the particles are put into cells at increasing trap potential until the trap is filled up with the available particles. At zero temperature there are no fluctuations and when one exceeds the chemical potential the cells remain empty, and there is a sharp $(d - 1)$ -dimensional interface between the occupied- and unoccupied trap states. This is of course the way we explain the Fermi-gas to our undergraduate students. It invokes an odd metaphor that however turns out to express an exact identification since we learned to handle the Ceperley path integral!

Having a statistical physics interpretation, can we now address the questions posed in section II? First, what is the order parameter of the Fermi-liquid? The answer is: the same order parameter that governs the Mott-insulator. This order parameter is well understood, although it is of an unconventional kind: it is the ‘stay at home’ emergent $U(1)$ gauge symmetry [76], stating that at every site and at all times there is precisely one particle per site. The particle number is locally conserved and henceforth a local $U(1)$ symmetry emerges. The ‘disorder operators’ that govern the finite temperature fate of the order parameter are just substitutional-interstitial defects: there is a finite thermal probability to excite a particle out of the trap, and the presence of the vacancies destroys the $U(1)$ gauge symmetry. Since the disorder operators are zero-dimensional particles regardless the dimensionality of momentum space, thermal melting of the Mott-insulator occurs at any finite temperature regardless dimensionality.

We repeat, this is just a rephrasing of the standard Fermi gas wisdoms in a non-standard language. The strange powers of the Ceperley path integral become more obvious when interactions are switched on. In the presence of the interactions single-particle momentum is no longer conserved, and this means that the worldlines of the Ceperley particles in momentum states get quantized: it is analogous to making the optical potential barriers finite in the cold gas Mott-insulator with the effect that the particles acquire a finite tunneling rate between the potential wells. One gets directly a hint regarding the stability of the Fermi-liquid: Mott-insulators are stable states that need a rather large tunneling rate to get destroyed. But the story is quite a bit more interesting than that, as can be easily argued from the knowledge on the canonical side. Let’s consider first what would happen in a literal cold atom Mott insulator when we start to quantize the atoms. Deep inside the trap motions are only possibly by doubly occupying the nodal cells and given that in the non-interacting limit the ‘Hubbard U ’ is infinite (expressing the Pauli surface) such processes are strongly suppressed. In the bulk of the trap the Mott state would be very robust. However, at the boundary one

can make cheap particle-hole excitations, and at any finite t the interface would no longer be infinitely sharp on the microscopic scale: the density profile would change smoothly. Eventually one would meet the ‘wedding cake’ situation where the bulk is still Mott-insulating while the interface would turn into a superfluid (we live in a bosonic world). How different is the Fermi-liquid! We know how it behaves from the canonical side. The single-fermion self-energy tells us directly about the fate of the \mathbf{k} -space Mott insulator. We learn that the time required to lose information on single-particle momentum is just given by the imaginary part of the self-energy and that behaves as [27] $1/\tau_k \sim (k - k_F)^2$, Henceforth, it diverges at the interface while it get shorter moving into the bulk. In the Ceperley bosonic language the Fermi-liquid is like a grilled marshmallow: It has a ‘crispy’, solid Mott insulating crust while it becomes increasingly fluid when one moves inside!

More precisely, the worldlines near the interface are fluctuating at short times, since we know that the momentum distribution of the bare electrons do smear around the Fermi-momentum - they do ‘spill out of the trap’. However, the effect of integrating out these microscopic fluctuations is to renormalize the ‘optical lattice potential’ upwards. This has to be the case because in the scaling limit the renormalized worldlines represent the quasiparticles and since they produce a perfectly sharp interface (i.e. unit jump in the quasiparticle n_k), the Mottness has to be perfect. This can only be caused by infinitely high effective potential barriers. This physics is of course coming from the modifications happening in the reach when interactions are turned on. The phase space restrictions giving rise to $\Sigma'' \sim \omega^2$ are rooted in Fermi-Dirac statistics and all the statistical effects are coded in the reach when dealing with the Ceperley formalism. These aspects can be computed by controlled perturbation theory and in a future publication they will be analyzed in detail.

3.4 The Fermi-liquid in real space: holographic duality

We showed in the previous section that at least for the Fermi gas the momentum space Ceperley path integral becomes a quite simple affair. Momentum space is a natural place to be when one is dealing with a quantum gas or -liquid, but dealing with a bosonic- or statistical physics systems one invariably runs into the general notion of duality [77, 78]. Dealing with conjugate degrees of freedom, like momentum and position or phase and number, one can reformulate the manifestly local order on one ‘side’ into some non-local topological order parameter on the dual side. An elementary example is the Bose-Einstein condensate. In the language of the previous section, one can either form a ‘black hole’ in the momentum space ‘trap’, by putting all bosons in the $\mathbf{k} = 0$ ‘optical lattice cell’. But one can also view it in real space, to discover the lively world of Section III where the local order in momentum space translates into a global, topological

description revolving around the infinite windings of worldlines around the time direction. Such duality structures are ubiquitous in Boltzmannian systems, and they are at the heart of our complete understanding of such systems: when one has a complete duality ‘map’ one understands the system from all possible sides and there is no room for surprises. For instance, when one is dealing with a strongly interacting system like ${}^4\text{He}$ one prefers the real space side because it is much easier to track the effects of the interactions [38]. Also in the strongly interacting fermion systems one expects that one is better off on the real space side. In this concluding section we will address the issue of the dual, real space description of the Fermi-liquid in the Ceperley path integral formalism. This real space side is remarkably complex: despite an intense effort even Ceperley and coworkers got stuck to the degree that they even did not manage to get things working by brute computer force. They ran into a rather mysterious ‘reference point glassification’ problem in their quantum Monte Carlo simulations, likely related to a contrived ‘energy landscape’ problem associated with the workings of the reach.

This is a fascinating problem: there has to be a simple, dual real space description of the Fermi gas. The obvious difficulty as compared to straightforward bosonic duality is the presence of the reach. One has to dualize not only the ‘life of the worldlines’ but also the constraints coding for the Fermi-Dirac statistics. Topology is at the heart of duality constructions and in this regard Ceperley [73], and more recently Mitas [79], have obtained some remarkably deep results, which will be discussed at length in the first subsection: the topology of the reach of the Fermi-liquid in $d \geq 2$ is such that the reach is open for all cycles of Ceperley worldlines based on even permutations or triple exchange. Henceforth, there is no topological principle that prevents infinitely long worldlines to occur and in subsection B we will argue that the zero temperature order of the Fermi-liquid has to be a Bose condensate of the ‘Ceperley particles’. This is conjectural but if it proves to be correct the Fermi-liquid holography we discussed in section 3.2 acquires a fascinating meaning: the scaling limit thermodynamics of the Fermi-gas in any spatial dimension $d > 1$ is governed entirely by the statistical physics associated with distributing the Ceperley worldlines over the cycles associated with even permutations, and this effective partition sum is indistinguishable from the partition sum enumerating the cycles of a soft-core boson system in one space dimension.

3.4.1 The topology of the Fermi-liquid nodal surface

To decipher the structure of constraints as needed for the real space Ceperley path integral one has to find out where the zero’s of the real space density matrix are. By continuation, these should be in qualitative regards the same in the Fermi-liquid as in the Fermi gas, and in the latter case we have an expression of the full density matrix in closed form,

$$\rho_F(\mathbf{R}_0, \mathbf{R}; \tau) = (4\pi\lambda\tau)^{-dN/2} \times \det \exp \left[-\frac{(\mathbf{r}_i - \mathbf{r}_{j0})^2}{4\lambda\tau} \right], \quad (3.13)$$

where $\lambda = \hbar^2/(2M)$. Henceforth, one needs to find out the zero's of this quantity for all \mathbf{R}_0, \mathbf{R} in the imaginary time interval $0 < \tau < \beta$. In real space, this is not an easy task. Part of the trouble is that at low temperature the zero's of the determinant depend on all coordinates at the same time. Only in the high temperature limit ($\tau \rightarrow 0$) the nodal surface of the density matrix becomes extremely simple [73]. To see this, define first a *permutation cell* $\Delta_{\mathcal{P}}(\mathbf{R}_0)$ as the set of points closer to $\mathcal{P}\mathbf{R}_0$ than to any other $\mathcal{P}'\mathbf{R}_0$. Obviously, the configuration space is divided into $N!$ permutation cells which are convex polyhedra bounded by hyperplanes, $\mathbf{R} \cdot (\mathcal{P}\mathbf{R}_0 - \mathcal{P}'\mathbf{R}_0) = 0$. The density matrix is simply a sum over all permutations and for $\mathbf{R} \in \Delta_{\mathcal{P}}(\mathbf{R}_0)$ and sufficiently high temperatures this sum is completely dominated by the term $(-1)^p \exp[-(\mathbf{R} - \mathcal{P}\mathbf{R}_0)^2/(4\lambda\tau)]$ since all the other terms are exponentially damped relative to it. Therefore, in the high temperature limit, $\rho_F(\mathbf{R}_0, \mathbf{R}; \tau)$ will have the sign of \mathcal{P} inside of $\Delta_{\mathcal{P}}(\mathbf{R}_0)$ and the nodal hypersurface is simply given by the common faces shared by permutation cells of different parities.

The reach acts both in a local way, much in the same way as we learned in the (1+1)-dimensional case as a special 'steric hindrance' structure having to do with entropic interactions, etcetera. However, it also carries global, topological properties and these are now well understood because of some remarkable results by Mitas [79], who managed to prove the 'two nodal cell' (or 'nodal domain') property of the higher dimensional Fermi-gas reach [73]. The topology of the nodal surface is associated with the structure of cycles as discussed in section III but now for the Ceperley path integral. The latter can be written as

$$Z = \sum_{\mathcal{P}_e} \int d\mathbf{R} \tilde{\rho}_D(\mathbf{R}, \mathcal{P}_e \mathbf{R}; \beta), \quad (3.14)$$

where \mathcal{P}_e refers to even permutations, while $\tilde{\rho}_D$ refers to the density matrix of distinguishable particles that are however still subjected to the reach constraints. As in the case of the Feynman path integral, this sum over even permutations can be recasted in a sum over cycles associated with all possible ways one can reconnect the worldlines at the temporal boundary, of course limiting this sum to those cycles that are associated with even permutations. We learned in section IV that for free worldlines even permutations translate into the supersymmetric quantum gas. But the Ceperley particles are not at all free, and the topology of the nodal surface tells us about global restrictions on the cycles that can contribute to Eq. (3.14).

It is immediately clear that the counting of cycles is governed by topology: to find out how to reconnect worldlines arriving at the temporal boundary from the imaginary time past, to worldlines that depart to the imaginary time future one needs obviously *global* data. This global information residing in the reach is

just the division of the reach in nodal cells we already encountered in the (1+1)-dimensional context and the momentum space Fermi gas. There we found that the space of all permutations got divided in $N!$ nodal cells, with the ramification that the sum in Eq. (3.14) is actually reduced to one cycles. Mitas has delivered the proof that in $d \geq 2$ the reach carries a two nodal cell topology, implying that all cycles based on even permutations lie within the reach. Since only this topological property of the reach can impose that certain cycles have to rigorously disappear from the cycle sum, this does imply that all cycles based on even permutations can contribute to the partition sum, including the cycles containing macroscopic winding numbers. Henceforth, the Ceperley worldlines can Bose condense in principle and it is now just matter of finding out what the distributions of the winding numbers are as function of temperature. This is what really matters for the main line of this story. Finding out the the way that Mitas determined the two-cell property is quite interesting and we will sketch it here for those who are interested. When you just want to understand the big picture, you might want to skip the remainder of this subsection.

Quite recently Mitas [79] proved a conjecture due to Ceperley [73], stating that the reach of the higher dimensional Fermi gas is ‘maximal’ in the sense that, for a given \mathbf{R}_0 and τ , the nodal surface of $\rho_F(\mathbf{R}_0, \mathbf{R}; \tau)$ separates the configuration space in just two nodal cells, corresponding with ρ_F being positive- and negative respectively. This is a quite remarkable property: for every pair \mathbf{R} and \mathbf{R}' in the same domain (lets say $\rho_F > 0$), one can change \mathbf{R} into \mathbf{R}' without encountering a zero crossing of ρ_F .

The easy way to prove this property goes as follows [79]. First, it can be demonstrated [73] that once there are only two nodal cells at some initial τ_0 than this property has to hold for any $\tau > \tau_0$. This follows straightforwardly from the imaginary time Bloch equation for the density matrix,

$$-\frac{\partial \rho(\mathbf{R}, \mathbf{R}'; \tau)}{\partial \tau} = H \rho(\mathbf{R}, \mathbf{R}'; \tau) \quad (3.15)$$

with initial condition,

$$\rho(\mathbf{R}, \mathbf{R}'; 0) = \det [\delta(\mathbf{r}_i - \mathbf{r}'_j)] \quad (3.16)$$

and the Bloch equation is a linear equation. This is a very powerful result because it gives away that the two-cell property ‘descends for the ultraviolet’: one has just to prove it at an arbitrary short imaginary time which is the same as arbitrary high temperature. Ignoring Planck scale uncertainties, etcetera, the form Eq. (3.13) has to become asymptotically exact for sufficiently small β , also in the presence of arbitrary interactions as long as they are not UV-singular! As we already noticed, this high temperature limit is rather tractable.

We now need to realize that we still have to take into account the ‘remnant’ of quantum statistics in the form of even permutations. Every even permutation can be written as a succession of exchanges of three particles $i, j, k \rightarrow j, k, i$

because these amount to two particle exchanges. When such an exchange does not cross a node (i.e. it resides inside the reach) the three particles are called ‘connected’. By successions of three particle exchanges one can build up clusters of connected particles. All one has now to demonstrate is that a point \mathbf{R}_t exists where *all* particles are connected in a single cluster, because this complete set of even permutations exhaust all permutations for a cell of one sign, because the odd permutations necessarily change the sign. One now needs a second property called tiling stating that when the particles are connected for the special point \mathbf{R}_t this has also to be the case for all points in the cell. And tiling is proved by Ceperley for non-degenerate ground states and also for finite temperature. Actually due to the linearity of the Bloch equation, its fixed node solution is unique, and the tiling property in the high temperature limit will lead to the same property at any lower temperature.

Before we prove that the above holds for the high temperature limit density matrix, let us just dwell for a second on what this means for the winding properties of the constrained path integral. The even permutation requirement means that, as for the standard worldline pathintegrals, we have to connect the worldlines with each other at the temporal boundary, but now we have to take care that we single out those cycles corresponding with even (or three particle) exchanges. The ‘maximal reach’ just means that cycles containing worldlines that wind an arbitrary large number of times around the time axis *never encounter a node* ! As noted before by Ceperley, this has the peculiar implication that in some non-obvious way the Fermi-gas has to know about Bose condensation. Since nodal constraints do allow for infinite windings there seems to be no ‘force in the universe’ that can forbid these infinite windings to happen and since the Ceperley path integral is probabilistic, when these infinite windings happen one has to accept it as Bose condensation. We will come back to this theme in a moment.

Following Mitas, one can now prove the two cell property of the high temperature limit using an inductive method. Assume that all N particles in the low β limit at a fixed \mathbf{R}_0 are connected in one cluster, to see what happens when an additional $N + 1$ particle is added. Single out two other particles $N - 1$, N and move these three particles away from the rest without crossing a node. Now we can profit from the fact that in the low β limit the density becomes factorizable: the determinant factors into a product of the determinant of the three special particles and the determinant of the rest. It is easy to show that the three particle determinant has the two cell property, proving that the $N+1$ ’s particle is in the cluster of N particles. Since this is true for any N , the starting assumption that all particles in the cluster is hereby proven.

For free fermions, Mitas also proved the two nodal cell property for non-degenerate ground states using a similar induction procedure. The trick is to choose a special point \mathbf{R}_t in the configuration space, at which one can easily show how all the particles are connected into a single cluster. Once proven for this single point, tiling ensures that the same is true for the entire nodal cell. Mitas

aligned the particles into lines and planes, thus forming some square lattice in the real space. This way the number of arguments of the wave functions is reduced and more importantly, the higher dimensional wave functions can be factorized into products of sine functions and the one dimensional wave functions, which are much easier to deal with than their higher dimensional counterparts. One distinct property of the 1 dimensional wave functions is that they are invariant under cyclic exchanges of odd numbers of particles, namely for N odd,

$$C_{+1}^x \Psi_{1D}(1, \dots, N) = \Psi_{1D}(1, \dots, N), \quad (3.17)$$

where C_{+1}^x represents the action to move every particle by one site in the $+x$ direction, with the last particle moved to the position of the first one, that is $1 \rightarrow 2, 2 \rightarrow 3, \dots, N \rightarrow 1$.

Consider for example the non-degenerate ground state of 5 particles in 2 dimensions. For this state, it becomes straightforward to show that each group of the 3 near neighbors living in the real space square lattice are connected by products of four triplet exchanges, which are all performed along the 1 dimensional lines. Proven this, one can proceed as in the high temperature limit, by adding more particles to the lattice. And these newly added particles can be shown to be connected to the original particles' cluster by the similar method used for 5 particles. The only difference is that now one needs to consider the whole line of particles, on which the new particle is added, and thus a sequence of four cyclic exchanges, instead of the special triplet exchanges are required. Since for non-degenerate ground states, there are odd number of particles on each line, cyclic exchanges will not produce extra minus signs, thus leading to the same result as triplet exchanges. This completes the proof for 2 dimensions, and the high dimensional cases are essentially the same.

However, winding is a topological property that should be independent of representation. In the long time $\beta \rightarrow \infty$ limit the path integral contains the same information as the ground state wave function, and for the Fermi-gas we can actually easily determine the winding properties inside one of the nodal cells using the random permutation theory. This demonstrates that at zero temperature the Fermi-gas is indeed precisely equivalent to the Bose gas, within the nodal cell.

3.4.2 There is only room for winding at the bottom

The conclusion of the previous subsection is that the Ceperley worldlines can in principle become infinitely long because the topology of the reach allows them to become macroscopic. Does this mean that the zero-temperature order parameter of the Fermi-liquid is just an algebraic bose condensate of Ceperley worldlines characterized by a domination of the partition sum by macroscopic cycles? The two nodal cell topological property is a necessary but insufficient condition for this to be true. However, there are more reasons to believe that the Fermi-liquid has to be of this kind.

However the zero- and finite temperature Fermi-liquid are separated by a phase transition and it appears that only the winding sector of the Ceperley path integral can be responsible for this transition. The argument is simple and general. With regard to ordering dynamics the real space Ceperley path integral is governed by Boltzmannian principle and let us find out what ‘substance’ is available to form an order parameter. The nodal surface in isolation cannot be responsible, since it is an immaterial object that just governs the behavior of the ‘Ceperley particles’. Henceforth, whatever its (singular) properties, these have to be reflected in the behavior of the matter. In principle one can imagine subtle topological changes occurring in the nodal surface but in the previous subsection we found this not to be the case in the Fermi-gas. Henceforth, searching for the thermodynamic singularity we should keep our eyes on the worldlines and these should be subjected to the generalities associated with bosonic matter. One source of thermodynamic singularity is that the system of bosons breaks the translational- and/or rotational symmetry of space, forming a crystal or some liquid crystal. Although the one dimensional Fermi-gas is such a crystal in disguise, it is impossible to hide a (partial) crystallization in higher dimensions: the higher dimensional Fermi-liquid is undoubtedly a true liquid. The worldlines have to be delocalized, but dealing with indistinguishable particles, being bosons or the ‘even permuting’ Ceperley particles, one has to account for an extra set of degrees of freedom: the reconnections at the temporal boundary. From a statistical physics perspective, Bose condensation appears as an order out of disorder phenomenon. Lowering temperature has the net effect of increasing the ‘configurational entropy’ associated with all possible ways of reconnecting worldlines, or either the appearances of cycles characterized by different windings. Worldlines get longer and thereby the length over which they can meander increases, and this in turn increases effectively the fugacity of long cycles. The more cycles can contribute, the larger the ‘configurational entropy’ associated with the cycles and this gain in space time ‘configurational entropy’ (physically the decrease of quantum zero point energy) causes eventually a flat distribution of the winding configurations, and in the Bose system this sets in at a sudden phase transition. Since all particles ‘are part of the same wordline’ the Bose condensate is macroscopically coherent. We learned that the reach allows the Ceperley particles to form infinite windings. We learn from the Bose condensate that at zero temperature only crystallization can prohibit the ‘reconnection entropy’ to take over, because the thermal de Broglie wavelength diverges. Henceforth, there does not seem to be any feature of the reach that can prohibit this to happen as well to the Ceperley worldlines at zero temperature.

There is a quite direct argument to support this view which was put forward by Ceperley some time ago [73,74]. As we already emphasized a number of times, on the canonical side the Fermi-liquid order manifests itself through the jump in the momentum distribution. Let us now turn to the zero temperature single

particle density matrix,

$$\begin{aligned} n(\mathbf{r}) &= \int d\mathbf{R} \rho(\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N; \mathbf{r}_1 + \mathbf{r}, \mathbf{r}_2, \dots, \mathbf{r}_N; \infty) \\ &= \int d\mathbf{k} e^{i\mathbf{k}\cdot\mathbf{r}} n_{\mathbf{k}}. \end{aligned} \quad (3.18)$$

In the boson condensate $n_B(\mathbf{r}) \rightarrow \text{constant}$ revealing the off-diagonal long-range order which is equivalent to the domination of infinite cycles. In the Fermi-liquid on the other hand,

$$n_F(\mathbf{r}) \simeq \frac{1}{(k_F r)^{d/2}} J_{d/2}(k_F r). \quad (3.19)$$

The oscillations governed by the Bessel function $J_{d/2}(k_F r)$ can be easily traced back to the size of the nodal pocket as discussed in a moment. However, the envelope function $(k_F r)^{-d/2}$ just behaves like the one particle density matrix of a Bose condensate showing off-diagonal long range order, like in the interacting Bose system in 1+1D at zero temperature. Relating this to the real space Ceperley path integral, this signals the presence of infinite cycles formed from Ceperley world lines.

