



Universiteit
Leiden
The Netherlands

Translational symmetry breaking in holographic strange metals

Balm, F.A.

Citation

Balm, F. A. (2023, May 16). *Translational symmetry breaking in holographic strange metals. Casimir PhD Series*. Delft-Leiden. Retrieved from <https://hdl.handle.net/1887/3618303>

Version: 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/1887/3618303>

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

I. Theory of Condensed Matter

In the early 1900's, the advent of the quantum revolution shone a new light on many of the open problems of physics at the time. An especially relevant example for this thesis is the conduction of electrons in metals. With the discovery of the electron at the end of the 19th century, Drude proposed a simple model for the conduction of electrons through solids. In this model, resistivity originates from a classical ideal gas of electrons scattering off each other like pinballs[1, 2]. The success of this theory was remarkable: both DC and AC electrical conductivity at room temperature are well described. The failings of this theory on the other hand provided a great challenge to the physics community. For example, transport properties at low temperature as well as heat capacities do not agree with experimental results at all. The eventual resolution of these problems have earned it its place as a starting point for modern physics education, as exemplified by its status as the subject of the very first chapter of esteemed textbooks like Ashcroft & Mermin[3]. In itself, it serves as an example of how the theory condensed matter physics has evolved over the past century. In relation to this thesis it will form a convenient point of departure from which to develop some of the ideas and challenges faced in condensed matter physics today, and how the the work contained in this thesis attempts to tackle them. In particular, I will present the case for the use of the holographic duality to study condensed matter systems, since some recent remarkable results from our research should attract the attention of the theoretical physics community at large to investigate classes of problems that are largely out of reach of conventional condensed matter physics theory.

I.1. Drude Transport

The Drude model of transport is a very straightforward textbook system, with which any physics student should be familiar. The full theory, as well as the subtleties related to the Sommerfeld improvements to the model, have been described in detail in many excellent resources[3]. The Drude model departs from the assumption that in a piece of metal the current density of electrons \vec{J} flowing through it can be expressed by

$$\vec{J}(t) = n e \vec{v}(t), \quad (1.1)$$

where n is the number density of carriers, e is the charge per carrier and \vec{v} is the velocity of each of the carriers. Assume now that an electron is scattered at a time $t = 0$ somewhere in the metal under the influence of a constant uniform electric field \vec{E} . At that point in time, it acquires some randomized velocity \vec{v}_0 due to the collision. After the collision, the electric field acts with a constant force $\vec{F} = e\vec{E}$ on the electron, meaning that its velocity at a time t after the collision is given

by

$$\vec{v}(t) = \int_0^t \frac{\vec{F}(t')}{m} dt' = \vec{v}_0 + \frac{e\vec{E}t}{m}. \quad (1.2)$$

In order to get the macroscopic behaviour, we have to perform an average over all particles in the system. It is assumed that this initial velocity \vec{v}_0 is uniformly randomly distributed such that the average value vanishes

$$\langle \vec{v}_0 \rangle = 0. \quad (1.3)$$

Hence this term drops out when taking the average

$$\langle \vec{v}(t) \rangle = \frac{e\vec{E}\tau}{m} \quad (1.4)$$

The relaxation time τ here is defined as the average time since a collision $\langle t \rangle \equiv \tau$. The classical result for a constant (DC) electric field is then that

$$\vec{J} = \sigma_{DC} \vec{E} = \frac{ne^2\tau}{m} \vec{E}, \quad (1.5)$$

where σ_{DC} is the DC conductivity. At room temperature this prediction works well when putting in estimates for number density and electron charge and mass.




Figure 1.1.: The Drude model departs from considering classical pinballs moving randomly under the influence of an electric field.

We can generalise the applied electric field to a general time-dependent force. In this case, we first express the electrical current \vec{J} in terms of the momentum \vec{P} of the electrons by

$$\vec{J}(t) = \frac{ne\vec{P}(t)}{m}. \quad (1.6)$$

Under the assumption that the electrons have some average scattering rate τ and are subject to some average force $\vec{F}(t)$ due to oscillating fields, by considering the probability of scattering in a

small time window dt and the effect of the force during that time, the differential equation governing the average momentum of the electrons can be written as

$$\frac{d\vec{P}(t)}{dt} = -\frac{\vec{P}(t)}{\tau} + \vec{F}(t). \quad (1.7)$$

Important here is that in the absence of $\vec{F}(t)$, we get a damping of $P(t)$, as the solution will go like

$$\vec{P}(t) = \vec{P}(0)e^{-t/\tau}. \quad (1.8)$$

In other words, if the electrons are not being driven, their average momentum will decay exponentially in time. The rate of the momentum relaxation will be intricately linked to the conductivity. Assume for example that the field is being driven at a fixed frequency ω , and that we are in a late-time system where the momentum also oscillates with this same frequency. Going to frequency space in equation (1.7), we can deduce that the frequency-dependent conductivity is given by

$$\sigma(\omega) = \frac{ne^2\tau}{m} \frac{1}{1 - i\omega\tau}. \quad (1.9)$$

The real part of this forms a Lorentzian in frequency space centred at $\omega = 0$ with width $1/\tau$.

In general, the presence of a Drude-like conductivity requires there to be a mechanism for momentum to decay. If it does not decay, it would mean that it lives forever, and the relaxation time would become infinite. As a result the conductivity would be infinite: all particles keep moving at ever increasing speeds, forever accelerated by the electric field, as they are never scattered. That would for example be the situation in the limit of a dilute Fermi gas in the Galilean continuum. In the real world though, there are many mechanisms through which momentum can decay. This is for example possible through Umklapp scattering, where an electron can dump a specific amount of momentum into the lattice, as well as through interaction with disorder in the lattice. Overall linear momentum of the whole system is still conserved, but appears to decay in the purely electronic part of the system. This is illustrated in figure 1.2.

1.1.1. Thermopower

As mentioned above, there are some areas where Drude theory does not work well, the thermopower is one of these. Consider what happens when a temperature gradient is applied to an electronic system. From elementary thermal physics, we know that electrons start moving from the hot side towards the cold side as a result along the direction of heat flow. This creates an electric field, as electrons are removed on the hot side and are transported to the cold side. We can define the thermopower Q as the ratio of induced electric field to temperature gradient [3]

$$\vec{E} = Q\vec{\nabla}T. \quad (1.10)$$

If we assume for now to be in a one-dimensional wire with spatial coordinate x , we can find that the average thermal velocity \vec{v}_Q is given by [3]

$$v_Q = \frac{1}{2} (v(x - v\tau) - v(x + v\tau)), \quad (1.11)$$




Figure 1.2.: Simple model showing Umklapp scattering for phonons. Through some process two phonons with momenta \vec{k}_1, \vec{k}_2 scatter to form a third phonon with $\vec{k}_1 + \vec{k}_2$, but as this is outside the first BZ this can be folded back by a lattice vector \vec{G} , such that $\vec{k}_f \neq \vec{k}_1 + \vec{k}_2$, meaning momentum gets lost.

or in other words, the average velocity that is induced by the temperature is given by the difference in velocities of electrons to the left and right of the point x . If there is a temperature gradient $\vec{\nabla}T$, these two velocities will be different, with the side with higher temperature having a greater velocity than the side of lower temperature. The factor $1/2$ comes in because roughly $1/2$ of the electrons will be moving in the $+x$ direction and the other half in the $-x$ direction. In a three-dimensional material, this will essentially apply equally in all three dimensions, where after some grinding of equations we can find that

$$\vec{v}_Q = \frac{\tau}{6} \frac{d(v^2)}{dT} \vec{\nabla}T. \quad (1.12)$$

From equation (1.4) above, we know that the electrons acquire an average velocity from the DC electric field which is given by

$$\vec{v}_E = -\frac{e\vec{E}\tau}{m}. \quad (1.13)$$

In a steady state the overall velocity which is a sum of the thermal and electrical velocity must be exactly zero, which means that

$$\vec{v}_E + \vec{v}_Q = 0. \quad (1.14)$$

Using the formulas above, the Drude model then predicts that Q is given by

$$Q = -\frac{c_v}{3ne}, \quad (1.15)$$

where c_v is the heat capacity at constant volume. This seems like a reasonable classical estimate, and one can wonder how this compares to real world physics. This is where the issues start: when one takes the classical kinetic theory result that $c_v = 3nk_B/2$, this goes enormously wrong: the value of the thermopower is overestimated by a factor 100 in electronic materials. The Drude model as it is presented here offers a clear picture of a classical view of conductivity, where resistivity can be phrased in terms of quantities like the collision time and the number density of charge carriers. As hinted at by the introduction, it is quantum physics that comes to the rescue here, and the resolutions of some of these issues are in a way a precursor to the issues that are faced in condensed matter physics today.

1.2. Beyond Classical Theory

With the knowledge of today, one can make the educated guess that it is the classical assumptions made in describing the electrons that are at the root of this problem. Instead, one should consider the electrons to be identical quantum mechanical particles. Furthermore, it is rash to ignore all the other interactions that are present in a real-world system, such as coulomb interactions between the conduction electrons for example.

1.2.1. A Better Electron Gas

In the time that Drude wrote down his theory, there was no clear picture of what had to be used as the distribution functions in the metals he was aiming to describe. Following arguments from kinetic theory, assuming an ideal classical gas of electrons would give rise to the Maxwell-Boltzmann distribution, which for a gas of particles moving in one dimension can be written as

$$f(v)dv = \sqrt{\frac{m}{2\pi k_B T}} e^{-mv^2/2k_B T}, \quad (1.16)$$

where k_B is Boltzmann's constant. At a fixed temperature, it is clear to see that the distribution has an $e^{-\alpha v^2}$ behaviour, which supposes that the vast majority of velocities drawn from this distribution will lie between $v^2 = 0$ and $v^2 = 4k_B T/m$. The shape of this distribution function has important consequences for some of the properties that we can compute about theoretical metals. In particular, as we saw before, when Drude calculated the thermopower, the specific heat that is found with this distribution is wrong by a very large factor. But if you suppose it must just be some wrong factor in the thermal transport, seeing that the electrical transport worked well, then it is all the more surprising that the Lorentz ratio between thermal and electrical conductivities as found in the Wiedemann-Franz law [4] matched rather closely to experimental results. This turns out to be a fortuitous coincidence, as we will soon see.

With the advent of quantum theory, it was realised that the Pauli exclusion principle must be applied to the electrons we are concerned with, as the underlying statistics that govern fermionic particles are antisymmetric in nature. This has far-reaching consequences: we cannot have multiple fermionic particles occupying the same state. Even in the ground state there are many fermions pushed to higher energy modes. Any undergraduate will have seen how this leads to characteristic

quantities that belong to a free electron gas, such as the Fermi energy and momentum and further derived quantities such as Fermi velocity and temperature. I will just highlight the Fermi velocity here: in a metal it turns out to be of the order of 1% the speed of light even at very low temperatures. This is in stark contrast to what is expected for the Maxwell-Boltzmann result, as that finds that the average thermal velocity would approach zero as the temperature approaches zero. At room temperature, the Fermi velocity is around two orders of magnitude greater than the velocity that Drude assumed in his original calculations.

For problems involving fermionic particles one should therefore not make use of the Maxwell-Boltzmann, but rather the Fermi-Dirac distribution

$$f(v)dv \propto \frac{1}{e^{\frac{\epsilon(v)-\mu_0}{k_B T}} + 1}, \quad (1.17)$$

where ϵ is the energy and v is the speed of the electron. This distribution has distinctly different features, especially the relationship with between v and μ_0 . For $\epsilon(v) \ll \mu_0$, this exponential is negligibly small and $f(v)dv \approx 1$. On the other side, when $\epsilon(v) \gg \mu_0$, the value of the exponential is large and therefore $f(v)dv$ is exponentially suppressed. The width of this transition is set by the temperature, T . For low temperature it is very abrupt, and for high temperature it becomes much more gradual. This is shown in figure 1.3. However, because the total integrated number of states that is described by the distribution must be the same at both low and high temperature, it must be spread out up to much higher velocities to account for all possible states.




Figure 1.3.: Difference between the low and high temperature behaviour of the Fermi-Dirac distribution. At high temperature a significant number of states above the chemical potential are occupied, whereas in the zero-temperature limit the transition becomes a step function.

Let us now revisit the matter of the thermopower as a good example, making only the simple change of going from the Maxwell-Boltzmann distribution to the Fermi-Dirac distribution.¹The

resulting expression for the thermopower is now augmented by an extra factor $\frac{k_B T}{E_F}$, which in many ordinary metals is of order 10^{-2} , which accounts for the factor that was missing from the Drude calculation from before.[3] Regarding the Lorentz ratio, it turns out there was another factor 10^{-2} overestimation, namely in the specific heat, indeed coincidentally cancelling each other out.

1.2.2. Towards Field Theory

This simple, semiclassical fix is a step in the right direction. One can wonder what happens whether a full, honestly quantum-mechanical treatment of a free electron gas would be able to capture all of the physics that is present in real metals. This is of course not the case: the interactions that are neglected in the free electron gas play an important role. Even more so, at no point in this discussion has the ionic lattice that makes up the metal itself been taken into account. This adds even more complexity, not to mention the interplay between the dynamics of the two.

One of the first steps further is to go beyond the world of simple quantum mechanics, and enter into a realm where it is possible to address more than just single particles. For example, if you wish to write own a wave function in quantum mechanics to deal with more than one particle that obeys fermionic statistics, you have to deal with fully anti-symmetrized wave functions. For a small number of particles, say (a, b) in positions $(1, 2)$, this is still tractable:

$$|\Psi\rangle = \frac{1}{\sqrt{2}} \left(|a(1)\rangle |b(2)\rangle - |b(1)\rangle |a(2)\rangle \right). \quad (1.18)$$

Problems quickly arise when we need to deal with more particles: the number of terms will grow as 2^n particles. In a macroscopic piece of material there are on the order of $n \approx O(10^{23})$ atoms in the game. The resulting number of terms, $2^{10^{23}}$ is such an unimaginably huge number that it will consist of $O(10^{22})$ digits. In terms of computer bits, this corresponds to about as much digital data as exists in the entire world, and that is only to write down the number of states itself - let alone writing down any of the individual states. To say that it is not a number that can easily be handled is understating it. A further problem is that quantum mechanics is not often very well equipped to deal with a non-fixed number of particles. A different approach is clearly needed. The answer comes in terms of field theory, where instead of single-particle degrees of freedom, we are instead dealing with a continuum of degrees of freedom, where we can think of the elementary low-level excitations as ‘particles’. I will not discuss the beauty and many subtleties of (quantum) field theory here.[5]

¹This requires some justification. Modifying only the probability distribution is only a part of the story if you are to do a proper quantum treatment. Quantum mechanics itself needs to somehow be included: this is nothing more than ‘classical pinballs’ but with a different velocity distribution. This semiclassical approach does give some satisfying results, and is sometimes good enough to get a decent theoretical prediction. A more detailed discussion can be found in [3].

1.3. (Broken) Symmetry

One of the central tenets in the description of physical systems is the concept of symmetry and what happens when symmetries are broken. It is an immensely powerful principle, not merely simplifying some mathematical expressions, but allowing us to disentangle the physical characteristics of a system into distinct components that are governed by different principles, separated by their symmetries, and the (conserved) quantities that are associated with each of them.

1.3.1. Noether's Theorem

Noether's theorem is perhaps one of the most famous ways in which symmetry manifests itself. For a classical system, Noether's theorem states that continuous symmetries of the Lagrangian of a system correspond to conserved quantities. That is, if we can make a transformation in the system that leaves the Lagrangian unchanged, there must be some sort of conserved quantity associated with that transformation. This has some well-known results associated with it, such as time-translational invariance being responsible for conservation of energy, spatial translational invariance providing conservation of momentum, and rotational invariance corresponding to the conservation of angular momentum.[6]

In classical field theory, the idea is analogous, only instead of conserved quantities, we get conserved currents. If we consider a Lagrangian in 4 dimensions that does not explicitly depend on the space-time coordinates x^μ , then the symmetry is encoded in the conservation of the stress-energy tensor $T^{\mu\nu}$, which can be evaluated as

$$T_\mu{}^\nu = \left(\frac{\partial \mathcal{L}}{\partial (\partial_\nu \phi)} \right) \partial_\mu \phi - \delta_\mu^\nu \mathcal{L}. \quad (1.19)$$

The conservation of energy and momentum can then be expressed as

$$\partial_\nu T_\mu{}^\nu = 0. \quad (1.20)$$

For quantum field theories, the situation is again very similar, and the role of Noether's theorem is played by the Ward identities in the theory.[5]

1.3.2. Landau Theory and Symmetry Breaking

While it is important to be able to identify conserved quantities, it is when symmetries are broken that things get really interesting. This is clearly seen in Landau theory.[7] Without going into any microscopic details of any theory we are dealing with, not making any reference to what physics governs these microscopics, we will depart from a simple model with a free energy given some scalar order parameter φ . For the sake of simplicity, let us only take temperature as an external parameter.

$$\mathcal{F}(T, \varphi) = A(T)\varphi^2 + B(T)\varphi^4. \quad (1.21)$$

The coefficients A, B are the ones that now fully determine the behaviour of this system. The form of equation (1.21) dictates that there will always be a local extremum at $\mathcal{F}(0) = 0$. We will

always take $B > 0$ here, such that the free energy is at least bounded from below. The interesting part is fully dependent on A . Schematically, there are two distinct regimes for what this free energy may look like, see figure 1.4.




Figure 1.4.: The free energy for the two different scenarios for $A(T)$. For $A(T) < 0$, two minima are present away where $\varphi \neq 0$, while there is a local maximum at $\varphi = 0$.

If we depart from the phase where $A(T) > 0$, then $\varphi = 0$ is the minimum of the free energy and the system will have that as its ground state. Imagine now that lowering the temperature changes the sign of $A(T)$ to be negative. When $A(T) < 0$, there will always be some range of values of φ where

$$\mathcal{F}(\varphi \neq 0) < \mathcal{F}(0), \quad (1.22)$$

which means that a phase exists with a non-zero order parameter which is at a lower free energy compared to $\mathcal{F}(\varphi = 0)$. In this particular example, the minima exist at

$$\varphi_{\pm} = \pm \sqrt{\frac{-A}{2B}}. \quad (1.23)$$

Therefore, the system will undergo a phase transition to a new phase where the order parameter φ takes on the value that minimizes \mathcal{F} . But that is not the whole story. From the explicit form of the model, it is clear to see that there is a manifest symmetry under the transformation $\varphi \rightarrow -\varphi$. This means that there is a priori no preference for the system to end up the minimum at φ_+ over the one at φ_- and vice versa. In the end, the system will, by some process or another, pick one of the two spontaneously². This is the essence of spontaneous symmetry breaking: the free energy preserves in this case the \mathbb{Z}_2 symmetry of the order parameter $\varphi \rightarrow -\varphi$. Without imposing a preference of one over the other the system will nonetheless end up in a state that breaks this symmetry spontaneously.

²One could argue that no matter the system, at some point even Heisenberg's uncertainty principle will always bring the system out of the unstable maximum at $\varphi = 0$.

1.3.3. Continuous Spontaneous Symmetry Breaking

The example above is the prototypical example of discrete spontaneous symmetry breaking. The concept of spontaneous symmetry breaking becomes even more physically interesting if the symmetry that relates the equivalent states is continuous. The standard example for this is the complex order parameter ϕ , with a similar (real-valued) free energy to equation (1.21):

$$\mathcal{F}(\phi) = a|\phi|^2 + b|\phi|^4. \quad (1.24)$$

The system depends only on the magnitude of the scalar ϕ . Again we assume that the constant b in front of the highest-power term of ϕ is always positive. For $a > 0$ there is a stable minimum at $\phi = 0$, but for $a < 0$, the point at $\phi = 0$ becomes an unstable maximum, as illustrated in figure 1.5. We can take advantage of the symmetry and write the complex scalar in terms of its magnitude and phase

$$\phi = |\phi| \exp(i\theta). \quad (1.25)$$

We can see that the value of the free energy is independent of the phase θ of the scalar. In figure 1.5, this means that there is now a circle at finite magnitude of ϕ that minimizes the free energy, there is a continuous $U(1)$ symmetry between all these states.




Figure 1.5.: Mexican hat potential for $a = -2, b = 1$.

The theory will have to pick some particular value of the phase θ to settle in, for which there is no particular preferential value. Here Goldstone's theorem comes into play. [8] This theorem states that for every generator of a continuous symmetry that is spontaneously broken, an extra massless scalar particle is added to the spectrum of the theory, the so-called Goldstone boson. These Goldstone bosons pop up in many places, the most important for our purposes is in crystalline solids. One can regard there the existence of the crystal lattice as a spontaneous symmetry breaking of

translationally invariant, isotropic, homogeneous space-time. The Goldstone bosons that pop up there we know more commonly as the phonons of the lattice.

Instead of spontaneous symmetry breaking, a theory can also be subject to explicit symmetry breaking. For example, one can start with a free-particle Hamiltonian, which is translationally invariant, and then add a periodic potential term to it. This modifies the theory, but now it is not only the states of the theory themselves that will break some symmetry: it is the Hamiltonian itself that is no longer translationally invariant. There is not such a neat theorem relating explicit symmetry breaking to appearing particles like there is for spontaneous symmetry breaking, but they can conspire and give rise to different effects.[9] For example, a symmetry is both spontaneously and explicitly broken, the boson that is generated by the spontaneous symmetry breaking now acquires a finite mass.

1.4. Superconductivity

All of the ideas discussed before come together in the theory describing superconductors. In superfluids, a global $U(1)$ symmetry is spontaneously broken. Superconductors are in a way similar, but there it is the $U(1)$ gauge redundancy that is broken³ spontaneously. The first superconductors were discovered by Heike Kamerlingh Onnes at Leiden at the start of the twentieth century, and the subject and its related puzzles continue to bring new and interesting questions more than a century later.[10] At first superconductivity had been observed only at temperatures very close to zero Kelvin. Advances towards the end of the 20th century saw the advent of the high- T_c superconductors, which in the 1980's and 1990's pushed the maximum temperature of the superconducting transition up from $O(10)$ Kelvin into the much more accessible cooling range of liquid nitrogen through ceramic compounds such as Lanthanum-Strontium-Copper-Oxides (LSCO) and Bismuth-Strontium-Calcium-Copper-Oxides (BSCCO).[11–15]

1.4.1. BCS Theory

In the first decades after the discovery of superconductivity, as it was found to be present in more and more materials, a theoretical explanation for the phenomenon was still sorely lacking. The very first microscopic theory came due to Bardeen, Cooper and Schrieffer.[16] Their insight was that an attractive interaction between fermions on other sides of the Fermi surface, no matter how small, has the effect of pair formation into what are now known as Cooper pairs.⁴ All manner of methods have since been developed to understand the physics of BCS theory, and for a concise introduction I will defer to reference [17].

³Which is again subtle and different from symmetry breaking as discussed before, as this generates massive bosons out of the photons rather than introduce new massless degrees of freedom. Regardless, symmetry breaking is still a key player here.

⁴The nature of the attractive interaction is one that for many intents and purposes can be left somewhat undetermined for quite a long time when considering BCS theory. One of the more conventional options is to consider the (screened) coulomb interaction in the solid in the presence of phonons. This can then be treated in perturbation theory to get to an effective interaction, which can be attractive. There is also the method of canonical transformations, which is described in detail in [18].

Another interesting aspect of BCS theory is the ground state that is assumed by the theory. In terms of creation, this is often written as

$$|\Psi_{\text{BCS}}\rangle \propto \prod_k \left(u_k + v_k c_{k\uparrow}^\dagger c_{-k\downarrow}^\dagger \right) |0\rangle. \quad (1.26)$$

Here u_k, v_k are some complex parameters, whose magnitudes must add to unity. This is a surprisingly simple waveform —instead of a complicated mixture of creation and annihilation operators concerning many different pairings on the Fermi surface, it turns out that the BCS state in momentum space is just a simple product state of pairs of created fermions. The importance of this statement will become clear later.

1.4.2. Challenges in Superconductivity

The BCS theory of superconductivity briefly displayed above works well for some compounds. Eliashberg theory, which in some sense can be seen as an extension of BCS theory in which retardation effects and dynamical interactions have been taken into account, works especially well for e.g. elemental superconductors.[19–21] One of the fundamental points of BCS superconductors is that the only symmetry that is broken is the $U(1)$ gauge symmetry, where the superconductor globally picks an (unobservable) phase spontaneously. The resulting order parameter is therefore highly symmetric. This preservation of symmetry is far from universal. In many systems of superconductors, additional symmetries are observed to be broken, and the symmetry of the pairing interaction can take the form of p -, d - or even f -wave symmetries. BCS theory is simply not equipped to deal with these more complicated symmetry breaking patterns. For example in the case of cuprate superconductors, iron pnictides, and many others, BCS theory simply fails.[21–23] These superconductors where the pairing interaction is not of BCS type are therefore known in general as unconventional superconductors.

Another commonly⁵ appearing characteristic in these unconventional superconductors is that it appears the correlations that appear in them are strong, much stronger than might be expected. Cuprate superconductors, especially those that are doped with holes, display this characteristic.[24] Some of these cuprate compounds have transition temperatures that exceed 100 Kelvin —much more than could ever be expected from a BCS superconductor.[25] One of the biggest outstanding questions in condensed matter theory exactly concerns describing the high temperature superconductors, and finding out why the temperature of their transition is so high. Many attempts have been made, but ultimately none so far have passed all the tests required in order to be accepted as ‘the’ theory of high-temperature superconductivity.[26, 27] But in spite of the attention it has received, it is perhaps not the most burning question in the cuprates from a fundamental physics point of view.

1.4.3. Hole Doped Cuprates

In the past decades, more than 200 different cuprate superconducting compounds have been identified.[28] What they have in common is their general crystal structure, which consists of one or

⁵But not necessarily.

more copper-oxide planes stacked on top of one another separated by layers of other atoms, such as Bismuth, Yttrium, Strontium, Calcium and more. In their undoped state, many of them fall into the category of Mott Insulators. In a quick-and-dirty classical way, Mott Insulators can be thought of as a material where the electrons in the conduction band get stopped in a ‘traffic jam’ due to strong lattice forces. The real story is of course rather more complex, but for present purposes it will suffice, and indeed strong potentials are the name of the game in the cuprates.[29] A whole world of phases opens up when these materials are doped, for example by changing the oxygen content, or by substitution of elements. Both the electron-doped and the hole-doped regions of the phase diagram are host to a plethora of interesting physical phenomena going on; we will be focused here solely on the hole-doped side.[28, 30]

While the details of all the different non-stoichiometric compounds created in this way are of course different, the phase diagram that can be mapped out of them follows some very general trends. This phase diagram, which is presented schematically in figure 1.6, contains much more than just Mott Insulators. In fact, although it is ‘unconventional’, the d-wave superconductor at the bottom is perhaps one of the least mysterious parts of this diagram, and it is rather the rest of the diagram that deserves much more attention. First of all, the Mott Insulator forms an anti-ferromagnetic state at small hole and electron doping, persisting up to a doping of a few percent. The superconductivity is only present at rather high doping, and is only truly ‘high T_c ’ in a somewhat small doping range around its maximum, the point of optimal doping. On the left of the superconducting dome is located the pseudogap phase.[31] This phase is rather mysterious and its transition temperature is rather hard to pin down exactly.[32] What can be said about it is that many of the different experimental techniques that are available, starting with nuclear magnetic resonance spectroscopy, have now confirmed that parts of the Fermi surface appear to become gapped far above the superconducting transition temperature.[21, 33] On the very highly doped side of the phase diagram, the situation becomes somewhat boring, and the materials return to being ordinary metals.

Many of the most interesting puzzles in the cuprates come together right in the centre wedge of the phase diagram, where strange metal phase resides.

1.5. Strange Metals

The strange metal has earned its name because it is quite literally a metal that behaving in a strange way. This normal-state region of the cuprate superconductors is really anything but that, and has several properties that display this general strangeness. The first and perhaps most striking characteristic is found in the in-plane resistivity⁶, which near optimal doping above the superconducting transition scales linearly with temperature [36–38]

$$\rho_{ab} \sim \rho_0 + bT. \quad (1.27)$$

⁶The residual resistivity at zero temperature, indicated by ρ_0 is not generically zero, but there are some systems where it does appear to vanish.[35] Rather, it is the constant linear slope that marks the linear-in- T behaviour.

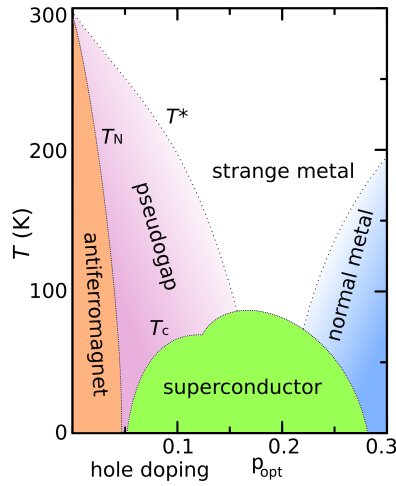


Figure 1.6.: Schematic phase diagram of cuprate superconductors on the hole-doped side. At zero doping the Mott insulator state is located. The pseudogap phase persists sometimes up to room temperature at low doping. When the doping is less than the optimal doping p_{opt} , this is referred to as underdoped, more doping is called overdoped. There is much that is not shown in this figure, such as possible charge density wave orderings and other phenomena.[34] T^* , T_c , T_N are the pseudogap, superconducting transition, and Néel temperatures respectively. This picture has been edited from https://commons.wikimedia.org/wiki/File:Cuprates_phasediagram_en.svg under the CC BY-SA 3.0 license.

This was seen already quite early on in the cuprate superconductors [39, 40].⁷ In some cases, this linear-in-temperature regime extends to far above room temperature. There are several reasons why this is interesting. First of all, the linear-in- T behaviour is not what is expected from a metallic phase: going by standard Landau Fermi Liquid theory, one expects the resistivity to instead scale as T^2 at low temperatures.[41] One typically expects the resistivity to saturate at a finite temperature, to what is called the Mott-Ioffe-Regel limit.[42] Another strange feature is that many of the cuprates blow right through that limit. The compound is still metallic, but conducts much worse than should be expected, hence this is often called ‘bad metallic’ behaviour.

This is only the start of the conundrum of the strange metal. This linear-in- T behaviour persisting over such a large range of temperature is an indication that there should be some very strong, universal principle underlying the physics of these systems. At the heart of this seems to be the hint that electronic interactions in the cuprates are strong, so strong that traditional perturbative methods are not able to capture the physics of these systems adequately. This is most sharply seen in the lack of quasiparticles, which seem to disappear in the strange metal, as well as the Hall angle,

⁷It is important to recognise that the ‘goodness’ of the linear-in- T resistivity does depend on doping, and it is very near optimal doping that the best linearity is often observed.[40]

does not behave as expected in a quasiparticle Fermi liquid.[43–47]

There is some evidence that right in the middle of the superconducting dome, at zero temperature, there is located a quantum critical point.[48] The scale invariance that is generated from the quantum phase transition, along with the suggestive characteristic wedge-shape of the strange metal region, seem to suggest that the strange metal might have some underlying quantum criticality. The nature of a possible quantum critical point is still a matter of debate in the physics community, as it is not sure what is actually happening at this point, and there is no good theory of what the order parameter might be that is undergoing the phase transition.[43, 49] The very high temperatures up to which the strange metal regime extends are also at odds with the idea of a single quantum critical point as the reason for all the strange behaviour.

Regardless of whether the quantum critical point actually exists, there are indications that there is some kind of quantum critical behaviour is going on. For example, measurements of the dynamic charge response $\chi''(q, \omega)$ seem to suggest that there is a special kind of quantum critical behaviour going on that completely decouples the spatial from the temporal dynamics.[50] The optical conductivity also shows interesting behaviour. Optical measurements show that there appears to be a very good Drude-like response at low frequency. One can describe a Drude response in terms of a width and a momentum relaxation rate. The Drude weight appears to be remarkably temperature-independent, but the Drude width appears to be set by some universal ‘Planckian’ time scale[51, 52]

$$\tau = A \frac{\hbar}{k_B T}, \quad (1.28)$$

where A is some $O(1)$ prefactor. In the typical Drude formula,

$$\sigma(\omega) = \frac{\omega_p^2 \tau}{1 - i\omega\tau}, \quad (1.29)$$

in the DC limit $\omega \rightarrow 0$, we recover that the resistivity goes as

$$\rho \sim \frac{1}{\omega_p^2 \tau}. \quad (1.30)$$

The Drude weight ω_p^2 , which is related to carrier density, is constant at low temperature and the relaxation time $\tau \sim T^{-1}$ (or equivalently, relaxation rate $\Gamma \sim T$). Together, this makes a very straightforward way to make the resistivity go linearly with temperature, $\rho \sim T$. This is just a set of observations though, and there is currently no good theory that unifies and explains all observations in a coherent way.

1.5.1. Strong Coupling Issues and the Fermion Sign Problem

The strange metal puts current state-of-the-art condensed matter theory on the spot. We require a theory where we do not think in that well-ingrained notion of quasiparticles that we as physicists are so familiar with, as the perturbative treatment that they are built on does not work in the presence of strong coupling. Common theories like the regular Fermi liquid, which works so well in normal metals, simply do not apply. The strange metal and in strongly correlated electron

systems are simply put obscured by their eponymous strong electronic couplings. If the couplings are weak, typically some kind of perturbation theory can be used, where the coupling strength λ is a small parameter. There are some ways of investigating strong coupling theories, and the strong interaction of hadrons has long provided a reason for high-energy theorists to try to explore it, but there is no good way of treating strongly coupled condensed matter systems in general. We will explain why.

So what can be done in order to describe the strange metal, if not analytically then at least numerically? One of the problems mentioned in section 1.2.2 was the size of the Hilbert space, which grows exponentially with the number of particles in the system. After all, Moore's Law predicts an exponential increase of our computing capabilities, and has been working well for decades now. One might hazard a guess that we are getting to a point of just letting the largest computers in the world loose on the problem and simply brute forcing an answer numerically.

Sadly, this is just a pipe dream. For example, the two-dimensional Hubbard model has often been proposed as the number one model that will with enough long-range hoppings be able to capture the relevant physics. In a Monte-Carlo setting at finite doping⁸, each of the configurations that needs to be considered carries with it an overall sign, depending on the exact realisation of the spins of the electrons in the system. The problem comes from the fact that to calculate an observable, one must average over a number of configurations. However, in order to get good statistics in a Monte-Carlo problem, observables need to be computed as an average of many configurations. The problem here is then that the number of configurations needed get to get acceptable error bars on the observables increases exponentially when going to lower temperature.[53] This growth is so fast, that it is simply impossible to make a large enough system.

Other approaches run into similar issues. Take for example the tensor networks, and DMRG in particular. [54, 55] Instead of having to deal with the entire Hilbert space, they rely on methods to somehow reduce the big $2^{10^{23}}$ Hilbert space down to the realm of what computers can handle. This is typically done by using some heuristic to decide on a number of states to keep. For some theories, this is easy. If the theory is gapped somewhere, it is safe to assume that the dynamics relevant to the ground state will not involve states above the gap scale.[54] This does not apply in the strange metal though. In the end, all numerical approaches always have to face that the number of states the Hilbert space needs to be reduced. In conventional, weakly coupled systems we can make some assumptions about entanglement for example, namely, that the ground state will take on some simple, product-state like form. This often works, and when thinking in terms of tensor networks this allows us to write the state as a product of complex matrices at different sites, each of which have some size χ called the bond dimension. It turns out the bond dimension can be interpreted to tell a lot about entanglement in the system. If a system has no entanglement at all, it is just a pure product of complex numbers (matrices of size $\chi = 1$) at the different sites, and it is only necessary to keep a minuscule part of the complete Hilbert space available for the lowest lying excited states. With increased entanglement, the size of the matrices that are needed to describe the entire state grows exponentially. Eventually, when the state is maximally entangled, we need again to take into account a large part of the Hilbert space. In the doped Hubbard model, it can be shown

⁸In the particle-hole symmetric case the problem is sign-free, however that is a well-known and well-studied case. This lucky simplification is not however true away from the exact symmetric point.

that one does need to keep a very large amount of states in the matrices in order to get reliable results - supporting the idea that dense entanglement can play a vital role in understanding the physics of this.[56] Therefore, it seems reasonable to want to be able to deal with densely entangled systems if one really wishes to describe the strange metal.

1.6. New Territory

The conclusions drawn from the previous section should all point in one direction. If we want to be able to describe the physics of the strange metal, we must deal with dense entanglement, strong correlations, strong lattice potentials, local quantum criticality and more, perhaps all at once, while not running into exponentially difficult computations.

From the unlikely realms of string theory and quantum gravity comes here a brilliant surprise: a duality that allows us to consider and compute theories at strong coupling, densely entangled, without running into fermion sign problems or exponentially large Hilbert spaces, where all the information is neatly encoded in black holes and curved space-times. This ‘holographic duality’ is at the very core of this thesis; it is using holography that all are computed and interpreted. In chapter 2, we will take a closer look at what this holographic duality encompasses and to what effect it has been used in this thesis. Chapter 3 will mainly focus on the definitions and conventions that are being used in the holographic lattices that are used throughout the thesis, and will outline the numerical procedures used to compute them.

The rest of the chapters will use and explore the holographic lattices in several different ways. Chapter 4 will deal with holographic probe fermions. By solving the Dirac equation in curved space-time, it is possible to compute the spectral function of strongly coupled fermions in the presence of very strong translational symmetry breaking. This indeed seems like it would have immediate connections to the properties of fermions in for example cuprate strange metals, where experiments like angle-resolved photoemission can measure this spectral function, however there are some important caveats, which will be discussed.

Chapter 5 is a rather technical chapter, but one that has played an important part, as it concerns a realisation that we had made that many groups that use the so-called Gubser-Rocha model misinterpret or overlook certain aspects of the model, especially in relation to the thermodynamical quantities that can be derived from this model. Furthermore, a correct interpretation of the model involving an exactly marginal scalar is convenient for use in the final chapter.

The most central paper of this thesis is chapter 6. In it, the Gubser-Rocha model for the holographic strange metal is subjected to both a 1D and 2D ionic lattices, with a wide range of lattice strengths. We then look at AC electrical and DC thermoelectrical conductivities to explore the precise mechanisms of transport in these lattices. Some highly intriguing results will be shown, including but not limited to a linear-in-temperature resistivity, the experimentally observed mid-infrared peaks in the optical conductivity, and a resistivity that appears to saturate to a relaxation rate set by a Planckian timescale.

Together, these chapters will form a succinct overview of uses of the holographic duality with an eye towards condensed matter physics and the strange metal.

