

A Thesis Submitted for the Degree of PhD at the University of Warwick

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/106449/>

Copyright and reuse:


This thesis is made available online and is protected by original copyright.

Please scroll down to view the document itself.

Please refer to the repository record for this item for information to help you to cite it.

Our policy information is available from the repository home page.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk



The metal-insulator transition in doped
semiconductors: an *ab initio* approach

by

Edoardo Giovanni Carnio

Thesis

Submitted to the University of Warwick for the degree of
Doctor of Philosophy

Department of Physics

May 2018

THE UNIVERSITY OF
WARWICK

Contents

Contents	i
List of Figures	v
List of Tables	vi
Acknowledgements	ix
Declarations	xi
Abstract	xiii
1 Introduction	1
1.1 Conduction in a doped semiconductor	2
1.2 Critical behaviour and universality	3
1.3 The exponent puzzle	4
1.4 Direct observation of Anderson localisation	7
1.4.1 Classical waves	8
1.4.2 Cold matter	9
1.5 The role of interactions	10
1.6 Approaches to include interactions	11
1.7 Organisation of this work	12
2 Silicon: properties and simulations	15
2.1 Electronic structure and DFT	16
2.1.1 Definition of the electronic structure problem	16
2.1.2 Solution via density functionals	17
2.1.3 The Kohn-Sham equations	19
2.1.4 Pseudopotentials	21
2.1.5 Linear scaling DFT	21
2.1.6 Forces and geometry relaxation	23

2.2	Fundamentals of silicon	24
2.2.1	Diamond cubic structure	24
2.2.2	Simulations of pure Si	24
2.2.3	Simulation of doped Si	26
2.3	Doping, shallow impurities and deep centres	29
2.3.1	Shallow impurities: doping with phosphorus	29
2.3.2	Deep centres: doping with sulphur	31
2.4	Defect pair formation and doping techniques	33
2.4.1	Growth of doped samples	34
2.4.2	Generation of disordered samples	34
3	From <i>ab initio</i> to effective models	37
3.1	The tight-binding model	37
3.2	Translation from DFT to effective models	40
3.3	Technical construction	43
3.3.1	Description in map	43
3.3.2	Creation of the catalogues	44
3.3.3	Construction of the effective matrices	45
3.4	Diagonalisation	47
3.5	Numerical challenges and the choice of sulphur	48
4	Multifractality of the wave functions	51
4.1	Continuous phase transitions	51
4.2	Multifractals	52
4.2.1	Measures, fractals and multifractals	53
4.2.2	The multifractal spectrum	55
4.2.3	Symmetry of the multifractal spectrum	56
4.3	Multifractal analysis of the wave function	56
4.3.1	Numerical calculation	57
4.3.2	Ensemble averaging	57
4.4	Statistics and disorder average	58
4.5	Results from the effective models	59
4.5.1	Negative fractal dimensions	60
4.5.2	Width of the multifractal spectrum	61
4.5.3	Symmetry of the multifractal spectrum	62
5	Finite-size scaling of the transition	67
5.1	Scaling theory of localisation	67
5.2	The problem with finite systems	69

5.3	Finite-size scaling of multifractal exponents	71
5.3.1	Derivation of scaling functions	71
5.3.2	Validity of the scaling assumption	72
5.4	Scaling at fixed coarse-graining ratio	74
5.4.1	Fits and scaling	75
5.4.2	An example	76
5.5	Other scaling possibilities	79
5.5.1	Inclusion of irrelevant corrections	79
5.5.2	Two-parameter scaling	81
5.6	Phase diagrams	83
5.6.1	Critical concentration	83
5.6.2	Critical exponent	83
5.6.3	Dependence on system size and resolution	84
5.6.4	Dependence on multifractal moments	87
6	Metal-insulator transition in the impurity band	89
6.1	Impurity conduction	89
6.2	Formation of the impurity band	90
6.3	Band gap closing and the metal-insulator transition	92
6.4	Delocalisation in the impurity band	94
6.5	Band merging and the critical exponent	96
7	Conclusions and outlook	99
A	Fit results	107
	Bibliography	115

List of Figures

1.1	Experimental estimates of the critical exponent in different materials throughout the years.	5
2.1	Essential blocks and keywords of the ONETEP input file to simulate pure Si.	25
2.2	Highest occupied Kohn-Sham eigenstate of 4096 Si atoms, as obtained from DFT.	27
2.3	Displacement vectors connecting a ‘1’ site to its first and second nearest neighbours in the diamond cubic lattice, expressed in units of lattice parameters.	28
2.4	Impurity state of a P impurity in the 4096 Si system, obtained from DFT	30
2.5	Impurity state of a S impurity in the 4096 Si system, obtained from DFT	32
2.6	Scheme of bands and impurity states in Si.	32
2.7	Distortion of the Si lattice induced by a P/S impurity.	33
2.8	Pair formation energy as a function of distance.	35
3.1	Example of localised and extended states.	39
3.2	Construction of the effective tight-binding model.	41
3.3	Comparison between 1- and 2-impurity catalogues.	42
3.4	Density of states of an exemplary system of 4067 Si atoms and 29 S impurities.	43
3.5	Example of a map of the doped Si system.	44
3.6	Dependence of the concentration on the number of impurities and atoms.	49
4.1	Multifractal spectrum for the ONETEP prototype described in Ch. 2.	56
4.2	Ensemble-averaged multifractal spectrum for a 22^3 system with 140 impurities.	63

4.3	Linear fits underlying the multifractal spectrum of Fig. 4.2.	64
4.4	Multifractal spectrum for the two mobility edges of a 22^3 system with 230 impurities.	65
5.1	Sketch of the asymptotic behaviour of the $\beta(q)$ function.	69
5.2	System-size invariance of $\text{PDF}(\alpha)$ at the Anderson transition.	73
5.3	Flow of $\text{PDF}(\alpha)$ under box-size renormalisation in the localised and delocalised regimes.	75
5.4	Example of scaling of α_0 for $\lambda = 1/2$	78
5.5	Example of α_0 data sets for $\lambda = 1/4$ and $1/6$	82
5.6	p -values for the fits with $\lambda = 1/6$	82
5.7	Phase diagram for $L \leq 22$, $\lambda \in \{1/4, 1/2\}$ and $q \in \{0, 1\}$	85
5.8	Phase diagram for $L \leq 20$ for $\lambda \in \{1/2, 1/4\}$ and $q \in \{0, 1\}$	86
5.9	Phase diagram for $L \leq 20$ and $L \leq 22$ for $\lambda = 1/2$ and $q = 0$	86
5.10	Phase diagram for $L \leq 20$ and $L \leq 22$ for $\lambda = 1/4$ and $q = 0$	87
5.11	Phase diagram for $L \leq 22$ and $q \in \{-1, 0, 1\}$	88
6.1	DOS of the IB for 4096 atoms at three different concentrations.	92
6.2	Band gap closing as a function of impurity concentration.	93
6.3	Shift of the Fermi energy as a function of impurity concentration.	94
6.4	Energy-localisation histogram for 4096 atoms with 40 and 100 impurities.	95

List of Tables

4.1	Summary of the ensemble statistics.	60
5.1	Test of the <i>robustness</i> of the fit described in Sec. 5.4.2.	77
5.2	Test of the <i>stability</i> of the fit described in Sec. 5.4.2	78
5.3	Results for the fits of the data in Fig. 5.4 ($\lambda = 1/2$) with an irrelevant term.	81
A.1	Fit results for $L \leq 20$ and $q = 0$	107
A.2	Fit results for $L \leq 20$ and $q = 1$	109

A.3	Fit results for $L \leq 22$ and $q = 0$	110
A.4	Fit results for $L \leq 22$ and $q = 1$	112
A.5	Fit results for $L \leq 22$ and $q = -1$	113

Acknowledgements

My sincere and deepest gratitude goes to Prof. Rudolf Römer: his role in my education goes beyond supervising me in this project, an admirable feat indeed. As a *Doktorvater*, he also led by example, teaching me invaluable lessons on professionalism, leadership, authenticity, and humanity.

I am also especially grateful to Dr. Nicholas Hine for his unfailing support with many aspects of our collaboration, and for bringing new ideas and a different perspective during our weekly discussions.

A heartfelt thank you goes to Prof. Andreas Buchleitner, for believing in me and guiding me in my development as a scientist and as a person.

These past years would not have been the same without the previous and current members of PS0.01, as well as the colleagues in and out the theory group: not everyone can count on a group of friends so supportive and caring. A special mention goes to the great moments with Álvaro, Anna (el Rey), Anto, and Edu: the funniest and most unproductive times of my PhD, and probably those where I have eaten more chocolate than ever. Siempre nos quedarán los Chupa Chups de Pamplona!

I also want to thank my friends in Freiburg, in particular Vanessa, Tobi, Luisa and Antonio, for always welcoming back with open arms. The same holds for my friends in Trieste: we may be spreading around the world, but the distance is only physical. I want to especially thank Arianna, Pietro, and Angelo for always opening their doors when I needed to get away and feel like at home for a weekend.

Finally, thank you to my “big sister” Sara and to my family, for always encouraging me to follow my dreams, even if that meant following me from far.

Declarations

I hereby declare that this Thesis has not been submitted in any previous application for a higher degree. This Thesis is submitted to the University of Warwick in support of my application for the degree of Doctor of Philosophy.

This Thesis represents my own work, except where references to other works are given. Most of Chapter 1, and portions of all other chapters, contain a scientific and historical background to the results presented: all provided information was gathered from the literature referenced either throughout the text or at the beginning of the corresponding sections. The methodology presented in Chapters 2 and 3, and the results of Chapters 4-6, are based on the following publication:

- E. G. Carnio, N. D. M. Hine, R. A. Römer, *Resolution of the “exponent puzzle” for the Anderson transition in doped semiconductors*, Under review (2017). Preprint available at [arXiv:1710.01742](https://arxiv.org/abs/1710.01742) [cond-mat.dis-nn].

The above-mentioned papers are the results of a joint work with the indicated authors.

Abstract

In this Thesis we study the Anderson metal-insulator transition starting from an atomistically correct *ab initio* description of a doped semiconductor. In particular, we use density functional theory to simulate model systems of sulphur-doped silicon (Si:S) with few impurities in a large cell. From the resulting Kohn-Sham Hamiltonian, we build an effective tight-binding Hamiltonian for larger systems with an arbitrary number of dopants. Our effective model assumes the same potential around single and paired impurities, for up to ten nearest neighbours and disregarding configurations of three and more close impurities. We generate up to a thousand disorder realisations for systems of 16^3 to 22^3 atoms and a large range of impurity concentrations. From the diagonalisation of these realisations we study the formation of an impurity band in the band gap of the host semiconductor. With increasing impurity concentration, this band undergoes an Anderson metal-insulator transition, namely (i) it approaches and merges with the conduction band and (ii) its states delocalise starting from the band centre. From the multifractal fluctuations of the wave functions near criticality, we characterise the Anderson transition in terms of its critical concentration n_c and exponent ν . We identify two regimes: for energies in a “hybridization region”, where the conduction band seems to influence the impurity band, we observe an increase from $\nu \approx 0.5$ to $\nu \approx 1$, compatibly with the experimental values; deeper in the band, instead, the estimates of ν fluctuate between 1 and 1.5, compatibly with $\nu \approx 1.59$ ($\nu \approx 1.3$) found in the Anderson model without (with) electron-electron interactions. Our results suggest a possible resolution of the long-standing exponent puzzle due to the interplay between conduction and impurity states.

Chapter 1

Introduction

We do not know if it was a dark and stormy night, but it was certainly the 1950s when Philip Warren Anderson was working on the unexpectedly long relaxation times of electron spins observed in the experiments on doped silicon conducted by Feher et al. (1955). Exactly sixty years ago this year, he proposed in Anderson (1958) a model where an electron diffuses via “quantum jumps” from site to site in a *random* lattice, and deduced that, when this randomness, also called *disorder*, is sufficiently strong, electron diffusion vanishes. Anderson’s idea was to place an electron on a specific site, let it evolve to the asymptotic time limit, and then compute the probability of finding it near the original site. When the disorder is stronger than a critical value, this probability remains *localised* around the starting position and decays exponentially over distance with a characteristic length ξ , called the *localisation length*.

Anderson localisation is fundamentally an interference effect, where a wave of wavelength λ is scattered elastically by the impurities of a disordered medium, with ℓ the mean free path between two scattering events. The comparison of these two length scales, together with the size of the system L , defines three different regimes (Feng and Jin, 2005). Firstly, if $\ell > L$, the wave propagates through an effectively homogeneous medium. When $\lambda < \ell < L$, instead, subsequent scattering events randomise the direction of the wave and transport is *diffusive*. It can be shown, e.g. in Sheng (2006), that, after multiple scattering, the backscattered waves interfere constructively. Since the increased backscattering opposes their propagation, the diffusion constant for waves must be reduced compared to the classical case. This effect, called *weak localisation*, is considered a precursor of Anderson localisation. When $\lambda/\ell \gtrsim 1$, called the Ioffe–Regel criterion (Ioffe and Regel, 1960), the wave

can resolve the inhomogeneity of the medium, the backscattering increases and the diffusion constant eventually vanishes, leading to *localisation*.

It took a number of years to recognise the broad repercussions of Anderson's work, after which a web of interconnected concepts developed way beyond the problem of impurity band conduction. Unsurprisingly, there are a number of excellent reviews that retrace this fascinating history, e.g. Al'tshuler and Lee (1988) and Lagendijk et al. (2009). In the next pages we will also expose the concepts we consider essential to contextualise our work.

1.1 Conduction in a doped semiconductor

What Anderson was seeking to describe is the diffusion of nuclear spins in a lattice or the conduction in the *impurity band* of a semiconductor (see Ch. 2). This band arises from the broadening of the impurity level that appears when we substitute one atom of the host material for a donor. In this context, disorder manifests itself as the randomly placed impurities perturbing the otherwise periodic potential landscape. Of the many themes connected to Anderson's work, then, we will review those related to the impurity band conduction problem, which is the topic of this Thesis. For this reason, it is time to introduce a second leading character of this story: Sir Neville Francis Mott, co-recipient with Anderson and John Van Vleck of the 1977 Nobel Prize in Physics. Anderson credits him for the progress and for keeping alive his interest in the twenty years after the paper on localisation (Anderson, 2010).

Mott had been working since the late 1940s on the metal-insulator transition (MIT), i.e. the transition from a metallic phase (overlapping bands) to an insulating phase (separated bands). If the disorder in the potential describing the impurity band of a doped semiconductor is strong enough, all wave functions are localised and transport is suppressed (Kohn, 1964). If it is not, Mott (1966) pointed out that localised states are more likely to appear first in the tails of the band. This is because electrons can hop only between states with similar energy, and at the tails there are fewer states available (Thouless, 1974). Localised and delocalised states should then be separated by an energy E_c , called the *mobility edge*. When the Fermi energy crosses E_c , e.g. by changing the doping, the system undergoes what Mott called an "Anderson transition". It took until the end of the 1980s for Mott and others to develop and test ideas on the emergence of conduction in an impurity band; we sketch some of this history in Chapter 6. The basic idea is

that, with increasing donor concentration, the impurity band of a semiconductor delocalises with an Anderson transition and merges with the conduction band of the host material.

1.2 Critical behaviour and universality

When the Fermi energy crosses the mobility edge into the localised phase, the conductivity σ of the sample continuously goes to zero. As we discuss in Chapter 5, this result stems from the work of Thouless, Wegner and then the “gang of 4” (Abrahams et al., 1979), who studied the Anderson transition from the perspective of critical phenomena. From this approach, called the *scaling theory of localisation*, two important concepts emerge: that the transition is universal and exists only for dimensions $d > 2$ (notice the strict inequality).

Universality is a hypothesis stating that most properties near a continuous phase transition do not depend on the microscopic details of the Hamiltonian, like the particular form of disorder or electronic and atomic interactions, but on *global* properties, such as its dimensionality and symmetries. Microscopic models could then be categorised in *universality classes*, identified by the values of the *critical exponents* that govern the power-law divergence of relevant thermodynamic quantities at the phase transition (Cardy, 1996).

As for dimensionality, the absence of a transition in one dimension had already been theorised by Landauer (1957, 1970), Anderson (1958) and Mott and Twose (1961). The interesting marginal case $d = 2$, instead, has been verified in numerical studies (MacKinnon and Kramer, 1981), as well as experiments on thin metallic films (Dolan and Osheroﬀ, 1979) and silicon MOSFETs (Bishop et al., 1980; Uren et al., 1981), but the prediction of the scaling theory holds for a gas of non-interacting electrons at the absolute zero and in the absence of magnetic fields. An Anderson transition in $d < 3$ could then occur when we consider a magnetic field, as we now discuss, or in the presence of interactions, which is the topic of Sec. 1.5.

In 1980, v. Klitzing et al. (1980) discovered that the Hall conductivity of a 2D electron gas in a transverse magnetic field shows discrete values, what is now known as integer quantum Hall effect (IQHE). These observations can only be explained by including the effects of disorder: because of imperfections in the system, the discrete states of a perfect crystal, called Landau levels, broaden into bands with tails of localised states, separated from delocalised states by a mobility

edge. The quantum Hall transition (QHT), the jump between discrete values of the conductivity, occurs when the Fermi energy moves between mobility gaps, where the states are localised, crossing a region of delocalised states. This clarifies why the study of the Anderson and quantum Hall transitions are strongly connected. For additional details, the reader can consult the reviews by Huckestein (1995), Kramer et al. (2005) and Evers and Mirlin (2008).

The other known possibility for an Anderson transition to occur in 2D is in the presence of spin-orbit coupling, as was first observed experimentally by Kravchenko et al. (1995). Compared to the standard Anderson model, the presence of a magnetic field or spin-orbit interaction breaks time-reversal T and spin-rotation S symmetries, respectively. Invoking universality, then, we conclude that these systems belong to different classes. When both symmetries are preserved, the Hamiltonian matrix is symmetric and thus invariant under rotations with orthogonal matrices (hence we speak of “orthogonal symmetry class”). On the other hand, when T is broken but S is preserved, the Hamiltonian is Hermitian and hence invariant under rotations with unitary matrices (“unitary symmetry class”). Finally, when S is broken but T is preserved, the Hamiltonian matrix is invariant under rotations with symplectic matrices (“symplectic symmetry class”). The topic is much broader and more complex than this short summary, and the reader is invited to start from Evers and Mirlin (2008) and references therein to find more details. As a summary, we want to point out here that, for $d = 3$, critical points exist for all three classes, while, for $d = 2$, they have been found only in the unitary and symplectic classes.

1.3 The exponent puzzle

Let us now focus again on the Anderson transition in the orthogonal symmetry class. In this case we are interested in the exponent ν describing the divergence of the localisation length with disorder, in the case of the Anderson model, or with dopant concentration, in experiments with doped semiconductors. The first of such experiments, by Paalanen et al. (1983) on phosphorous-doped silicon (Si:P), reported $\nu \approx 0.5$ on both sides of the transition. In the metallic regime, i.e. for concentrations n larger than a critical value n_c , the scaling of the conductivity as $\sigma \propto (n/n_c - 1)^\nu$ is measured. Notice that, at least in 3D, σ scales with the same exponent of the localisation length, as we show in Sec. 5.1. In the insulating regime, instead, the metal-insulator transition is observed via the scaling of the




Figure 1.1: Experimental estimates of the critical exponent in different materials throughout the years. We have collected the values from Thomas (1985), M. Itoh et al. (2004), and references therein. Whenever the original authors did not provide an estimate on the uncertainty, we have set it to 0.

donor polarizability $\chi \propto (n_c/n - 1)^{-2\nu}$ (Rosenbaum et al., 1980).

In the following years, many other experiments were performed, with various techniques and on different materials. Interestingly, they also found different exponents, as shown in Fig. 1.1. The results on different materials seem to fall in two categories based on the degree of compensation, namely the presence of acceptors in addition to donors. Reviewing the experimental results from those years, Thomas (1985) notices that for uncompensated materials $\nu \approx 0.5$, while for compensated semiconductors and amorphous materials $\nu \approx 1$. An explanation for the different behaviour in Si:P and other uncompensated semiconductors was proposed by Stupp et al. (1993), where $\nu \approx 1.3$ was found by (i) narrowing the critical region when performing the scaling analysis and (ii) including only the data where a reliable extrapolation to temperature $T \rightarrow 0$ is possible. Following these precautions, some years later Waffenschmidt et al. (1999) repeated the same experiments on Si:P performed by Paalanen et al. (1982) and found a critical exponent of $\nu \approx 1$ instead of 0.5. In the same year and in a similar experiment, Bogdanovich et al. (1999) found the critical exponent for uncompensated Si:B to be $\nu \approx 1.6$, a significant increase from the initial $\nu \approx 0.65$ found by Dai et al. (1991).

In this case, the discrepancy with the results of Waffenschmidt et al. (1999) might be attributed to higher temperatures, which do not allow a reliable extrapolation to $T \rightarrow 0$, or to strong spin-orbit scattering (Evers and Mirlin, 2008).

Further references and a summary of the different experimental estimates of ν can be found in M. Itoh et al. (2004). In this work, the authors use a novel doping technique, based on neutron-induced transmutation, to study the MIT in Ga:Ge. This method allows the production of homogeneously-doped samples, as opposed to melt-doping methods, which can yield significant fluctuations in the doping concentration between samples. The authors study the transition both with and without compensation and magnetic field. In all cases, the critical exponent is found to be around 1, with the caveat, for the uncompensated samples, that the critical region is restricted to 1% of the critical concentration of impurities. Outside of this region, the authors cannot scale their data with temperature; they therefore rely on the $T \rightarrow 0$ extrapolation used in previous studies and find again $\nu \approx 0.5$. A number of factors, then, can influence the estimate of the critical exponent: (i) the technique used to grow the samples, and the degree of homogeneity they can achieve throughout a crystal, (ii) the degree of purity, since perfect compensation is impossible to achieve in experiments, (iii) the breadth of the critical region, and (iv) temperature. As long as the critical region is kept narrow, however, the picture consistently points to a value of $\nu \approx 1$. Two points remain unexplained so far. The first is the connection between the degree of compensation and the critical region: the authors speculate that $\nu \approx 0.5$ is actually the *intrinsic* behaviour of the uncompensated semiconductor, although this would disagree with scaling theory and its consequences (e.g. $\nu \geq 2/3$ found by Harris (1974); Chayes et al. (1986)). The second is that the estimate of ν from theoretical models is *not* 1, as we now review.

After the publication of the scaling theory of localisation in 1979, MacKinnon and Kramer (1981) and Pichard and Sarma (1981) independently developed the *transfer matrix method* (TMM) to calculate the localization length ξ . As we discuss in Ch. 5, the scaling of ξ/L with linear size L and disorder W gave the estimate $\nu = 1.2 \pm 0.3$, which is compatible with the experimental values (MacKinnon and Kramer, 1981). Some years later, however, Schreiber et al. (1989) showed that, by increasing the accuracy of the TMM, the estimate of the exponent shifted to $\nu = 1.5 \pm 0.2$. From then on, further refinement of the calculations and of the scaling methods led to a reduction of the estimate uncertainty to $\nu = 1.57 \pm 0.02$ (Slevin and Ohtsuki, 1999). This estimate was confirmed by numerical simulations

with wave-function-based measures over large disorder ensembles, which resulted in $\nu = 1.590(1.579, 1.602)$, the most accurate estimate of ν in the non-interacting 3D Anderson model (Rodriguez et al., 2011).

It is interesting to note that a similar discrepancy appears also in the study of the quantum Hall transition (2D unitary class), where experiments find $\nu \approx 2.38$ (Li et al., 2009) and numerical simulations of the Chalker-Coddington model (Chalker and Coddington, 1988) find $\nu = 2.593 \pm 0.005$ (Slevin and Ohtsuki, 2009). For a review of the numerical estimates of ν of the Anderson transition in different symmetry classes and dimensions, the reader can consult Slevin and Ohtsuki (2014).

1.4 Direct observation of Anderson localisation

At the turn of the millennium it was clear that numerical simulations and experiments did not agree on the value of the critical exponent. A vast sea separates these two shores. On the one hand, it seems that the simplicity of the Anderson model, lacking the treatment of interactions or material-specific features (crystal structure, presence of nuclei, electronic configuration), is insufficient to capture the physics of the impurity band of a real material. On the other hand, doped semiconductors are too complex systems for experiments to single out any specific factor that drives or influences the localisation transition. In these systems, after all, we do not observe the explicit (de)localisation of the electron wave function, but rather its indirect consequences as a MIT, measured in the transport experiments described in the previous section.

The problem with a *direct* confirmation of localisation is that electrons are in fact quite “sticky”: they interact with each other and also with phonons. Hence, it becomes rather difficult to disentangle the role of disorder from that of interactions. Anderson localisation, however, is a universal *wave* phenomenon and should occur also in classical waves, as shown by Sajeed John (John et al., 1983; John, 1984). Shortly thereafter, Anderson himself proposed some experiments with elastic and electromagnetic waves (Anderson, 1985), the latter attracting a lot of interest due to the weakly-interacting nature of photons and the large coherence lengths available with lasers. The possibility of localising light could lead to significant technological advancements in several fields, from photovoltaic to medical devices (Segev et al., 2013; Wiersma, 2013).

1.4.1 Classical waves

Localisation of light has been quite elusive so far. From the theoretical point of view, localised states of light could be observed in a medium with high refractive index and at an intermediate frequency, that of Mie scattering, where the wavelength is comparable to the typical size of the scatterers (John, 1984). As for experiments, the first to claim the observation of three-dimensional light localisation was Wiersma et al. (1997), although it later turned out that these results could be described equally well by residual absorption (Scheffold et al., 1999). Ten years later, Störzer et al. (2006) reported localisation of light in an experiment with compressed titanium-oxide powders, a similar material to the “white paint” suggested by Anderson (1985) himself. By studying the time-resolved profile of short pulses transmitted through the samples, the authors can separate the different effects of disorder and absorption. It took ten additional years for the same team to ascertain that, in fact, their observations could result from fluorescent emission processes, and that Anderson localisation of light in 3D remains unobserved (Sperling et al., 2016). As remarked by Skipetrov and Page (2016), while the refractive index of common transparent materials might be insufficient for these observation, *ad hoc* engineered materials might provide a successful avenue. Photonic crystals, for instance, had already been proposed by John (1984) and have been employed to observed localisation of light in 2D (Schwartz et al., 2007).

Localisation of sound waves, instead, was first reported by Hu et al. (2008). In this experiment, a ultrasound pulse was transmitted through a cylindrical sample composed of aluminium beads forming a random network. The frequency of the pulse was 0.2–3 MHz, such that the wave length was comparable with the bead and pore size. The sound wave hit one of the faces in its centre and the transverse intensity profile $I(\rho, t)$ was measured on the other side, at radii ρ away from the centre $\rho = 0$. By studying $I(\rho, t)/I(0, t)$ any time-dependent absorption effect is cancelled. In the diffusive regime, the transverse intensity profile spreads in time like a Gaussian wave packet, while in the localised regime it saturates. A similar setup was later used in Faez et al. (2009) to present the first experimental observation of the multifractal nature of the “wave function” (in this case the normalised intensity profile) near the localisation transition (see Ch. 4). The same research group further improved their experimental investigation in Aubry et al. (2014), Hildebrand et al. (2014). In Cobus et al. (2016), in particular, they demonstrate the existence of an Anderson mobility gap and calculate the localisation length as a function of frequency, obtaining $\nu \sim 1$ as a first rough estimate of the critical

exponent.

1.4.2 Cold matter

In the same year, two experiments (Billy et al., 2008; Roati et al., 2008) observed the spatial localisation of a non-interacting Bose-Einstein condensate (BEC). Atomic gases are an excellent platform to study problems in condensed matter theory, owing to the control over atomic interactions and disordered potentials. Furthermore, the direct imaging of condensates allows the measurement of properties, like the localisation length, which were unavailable before.

In Billy et al. (2008), the BEC was confined in a one-dimensional wave guide using a weak magnetic field. The confining field was then switched off and the condensate was allowed to expand, driven by the initial repulsive interaction energy, in the presence of a disorder potential. The latter was generated using an optical speckle field, i.e. a random intensity pattern formed by the diffusion of coherent laser light on an irregular surface. This process yields a disorder potential with no long-range correlations and whose parameters can be tuned with high accuracy (Clément et al., 2006). In the presence of disorder the expansion of the BEC soon halts, and after 1 s the density profile along the wave guide shows exponentially localised tails.

Roati et al. (2008) used instead a BEC on a one-dimensional quasi periodic lattice to realise experimentally the Aubry-André model (Aubry and André, 1980), which shows an Anderson transition already in 1D. Two parameters can be independently tuned: the tunnelling between sites J and the disorder strength Δ . The time evolution of the spatial distribution of the condensate was studied using absorption imaging, and fitting with an exponential decay showed a smooth crossover between ballistic diffusion ($\Delta/J = 0$) and localisation ($\Delta/J \gg 1$).

A few years later, similar experiments with ultracold matter were conducted to observe Anderson localisation in three dimensions. In both cases, ultracold atoms – of fermionic nature for Kondov et al. (2011) and a BEC for Jendrzejewski et al. (2012) – were allowed to expand in a 3D optical disorder potential. In 3D we expect to see an Anderson transition and hence a mobility edge E_c . This was directly measured first by Semeghini et al. (2015) by developing techniques to control the energy distribution. In this experiment, a BEC is first loaded into the low-energy states below E_c ; a controlled perturbation of the disorder is then used to excite a fraction of these states above E_c ; these states are then allowed to

diffuse through the system and become invisible to the imaging device used in the experiment; the missing states are finally used to determine the position of E_c .

Experiments with cold matter are very exciting because they are a platform with controllable disorder and energy, and hence open the door to a laboratory realisation of the Anderson model. An experimental test of its universality has already been performed by Lopez et al. (2012) using a quasi-periodic atomic kicked rotator, where they found an average critical exponent of $\nu = 1.63 \pm 0.05$. Semeghini et al. (2015) also mention the possibility of measuring ν by narrowing the energy distribution in future experiments.

Two decades into the millennium, then, we can say that our experimental techniques have become so advanced that we can now replicate in the lab, without external disturbances, simplified and paradigmatic models of reality – the vision of a *quantum simulator* (Johnson et al., 2014). How have theoretical models fared in resolving the exponent puzzle?

1.5 The role of interactions

Progress in understanding the Anderson transition in the presence of interactions has been slower than its non-interacting counterpart. One of the first problems considered was whether electronic interactions would support *hopping conductivity*, the phonon-assisted hopping between localised states of an electron that is not excited to the mobility edge (Miller and Abrahams, 1960). Fleishman and Anderson (1980) showed that “short-range” interactions (including Coulomb for dimensions $d < 3$) are *not* sufficient to cause transport without phonons or delocalised states, hence the conductivity remains zero. This was later confirmed and further developed by Gornyi et al. (2005) and Basko et al. (2006).

Another question that soon arose is whether interactions are able to induce a MIT in two dimensions, as reported by many experiments on thin films performed in the 1990s (see Abrahams et al. (2001) for a review). The founding principles to answer this question were laid in the work of Finkel'shtein (1983), which extends the field-theoretical description of the Anderson model (Wegner, 1979) to include the effect of interactions. For the many subsequent theoretical developments, please consult Belitz and Kirkpatrick (1994). One of the most important results that followed was the work by Punnoose and Finkel'stein (2005), which presents a theory of the MIT in a disordered 2D electron gas. The phase transition emerges because interactions have an “anti-localisation” effect that contrasts the presence

of disorder and stabilises the metallic phase (Finkel'stein, 2010).

An important development in the study of the MIT in 2D materials comes from the work of Richardella et al. (2010), who used scanning tunnelling microscopy (STM) to visualise the electronic states of a two-dimensional dilute magnetic semiconductors. The correlation length, computed from the spatial structure of the states, shows a divergence at the Fermi energy when increasing the sample doping, the trademark of a MIT. At the same time, the multifractal spectrum (see Ch. 4) shows a characteristic shift towards the weak multifractality limit, identified in a reference delocalised state from the valence band. This work proves that the critical correlations in the electronic states near the MIT are not necessarily washed away by the presence of the Coulomb interaction (Burmistrov et al., 2013), but are even accessible in an experimental context. Like in experiments with cold matter, imaging techniques could prove very insightful in understanding how interactions affect localisation.

1.6 Approaches to include interactions

As pointed out in the perspective by Thouless (2010), electron-electron interactions can be included in theoretical models using the Hartree-Fock method or density functional theory (DFT, see Ch. 2) by Kohn and Sham (1965). In fact, the idea of using DFT to study the MIT in doped semiconductors is much older (Ghazali and Hugon, 1978).

In 2014, Amini et al. (2014) performed a Hartree-Fock simulation to study disordered, spinless fermions interacting via the Coulomb potential, whose intensity U is used to drive the transition. The authors confirm the persistence of multifractality and compute, from the divergence in U of the correlation (localisation) length ξ , two estimates for the critical exponent: $\nu_I \approx 1$ for the insulating regime, and $\nu_M \approx 0.5$ for the metallic regime. The scaling of ξ , however, is studied only for one system size, and fitting the metallic/insulating branches separately contradicts the scaling theory of localisation, which assumes that *one* single function describes both regimes.

Another numerical work reporting an estimate of the critical exponent when including the Coulomb interaction is Harashima and Slevin (2014). Here, the donors in Si:P are distributed randomly in a medium with the effective mass and dielectric constant of silicon. The Schrödinger equation is solved self-consistently using density functional theory (DFT), which offers an approximation to the

Coulomb interaction via an *exchange-correlation functional* (see Ch. 2 for an introduction) in the local-density approximation (LDA). The authors study the multifractal properties of the highest-occupied single-electron state and, similarly to previous studies by Slevin and Ohtsuki (1999), perform a finite-size scaling analysis of the transition (see chapters 4 and 5) yielding an estimate $\nu \approx 1.3$. The deviation in critical exponent from $\nu \approx 1.59$ of the non-interacting case is consistent with the idea that the Coulomb interaction should change the universality class of the transition (Burmistrov et al., 2013). Despite the intrinsic approximations of DFT, acknowledged by the authors, this work has the merit of bringing together density functional theory, which is routinely and reliably employed to study the properties of (often more complex) materials, with the finite-size scaling analysis of the transition, which proved so successful in estimating ν in the non-interacting Anderson model. We must nevertheless remark that the estimate of Harashima and Slevin (2014) is still *not* compatible with $\nu \approx 1$ found in experiments. Their model, after all, describes donors in an effective medium, which might be an oversimplified description of a doped semiconductor.

1.7 Organisation of this work

Considering the current advance in high-performance computing, we set out to study doped semiconductors *ab initio*. We use density functional theory to simulate a volume of silicon atoms, with a fraction of donor impurities. In its more modern incarnations, DFT is able to simulate up to 10^3 - 10^4 atoms on adequately-sized supercomputers, but these tasks usually require several hours on many hundreds or thousands of computing cores. We cannot realistically perform a disorder average to study the MIT. For this reason, in this Thesis we use the Kohn-Sham Hamiltonian to construct *effective* tight-binding models of doped semiconductors. We drive the MIT by changing the concentrations of impurities and the size of the systems. After computing the eigenstates of the tight-binding Hamiltonians, we perform a finite-size scaling analysis of the multifractal exponents to characterise the transition (Rodriguez et al., 2011; Harashima and Slevin, 2014).

The organisation of this thesis follows the stages of the work flow. **Chapter 2** is about the prototypical DFT simulations of doped silicon that we use afterwards to generate the effective models. Accordingly, we will review the basics of density functional theory, of silicon, and of doping, comparing them with the results of our simulations. Finally, we discuss how disorder is realised by randomly

distributing impurities in the silicon lattice.

In **Chapter 3** we move to the construction of the effective models. We briefly present the tight-binding model and connect it to both the Anderson model and the Kohn-Sham Hamiltonian. The discussion continues with the assumptions and technical details involved in reusing the DFT results to build the effective models. We touch upon the topic of diagonalising the tight-binding models and the limitations imposed by the available computing power.

From the diagonalisation we obtain eigenvectors and eigenvalues. The statistical properties of the former are analysed using multifractal analysis (MFA), which is the theme of **Chapter 4**. We give an introduction to multifractals and how they arise in the study of continuous phase transitions. As an example, we show the multifractality of the DFT-simulated wave function. We extend MFA to include the averaging over disorder and apply it to investigate the properties of our effective model.

Chapter 5 is wholly devoted to the finite-size scaling analysis of the transition. For this reason, we introduce the scaling theory of localisation and finite-size scaling analysis, followed by a discussion of the fitting procedure and the reliability of the scaling function we use. We show the phase diagram of the critical concentration and exponent as a function of the energy in the band gap, and test the robustness of these results with respect to the system sizes, the coarse graining applied to the wave functions, and the multifractal moment. Our data suggests that the impurity band delocalises first in the centre, and that the critical exponent increases from $\nu \approx 0.5$ at the Fermi energy to 1–1.5 deeper in the impurity band.

In **Chapter 6** we talk about the MIT in the impurity band (IB). We introduce the work of Mott on the conduction in the IB and identify the conditions for a MIT as the merging of the band and a delocalisation of its states near the Fermi energy. From the eigenvalues of the effective models we observe the formation of an impurity band, while the delocalisation transition was studied in Ch. 5. By combining the information on the merging with the conduction band, and the localisation of the states in the IB, we justify the trends observed in the phase diagrams.

Chapter 2

Silicon: properties and simulations

The first step of our workflow consists in simulating prototypes of doped silicon *ab initio*, and that is the focus of this chapter. We introduce the electronic structure problem and density functional theory (DFT) as a method to solve it. This thesis is about Anderson localisation, hence we include only the minimal information on DFT required to perform the simulations in this work: the Hohenberg-Kohn theorems and the role of the electron density distribution, the Kohn-Sham equations, pseudopotentials, linear-scaling DFT and nonorthogonal bases, calculation of forces, geometry relaxation.

We then discuss how to simulate a large cube of solid silicon, as the results of these calculations will be the basis of the chapters that follow. We focus on the crystal structure and its symmetries, as this will help us decide how many prototypes we need to study when we consider single and pairs of impurities.

We also introduce basic concepts in the electronic properties of defects, comparing shallow impurities and deep centres, exemplified, respectively, by phosphorus and sulphur. In this summary we already show and compare the results from some of our own DFT calculations. We conclude with some observations on the creation of pairs of defects, defining an “effective potential” between impurities, on the techniques to grow samples of doped silicon, and on how these are reflected in the generation of disordered samples to study with the effective models introduced in the next chapter.

2.1 Electronic structure and density functional theory

In this section we discuss the electronic structure problem and how density functional theory (DFT) can be used to solve it. We follow the logical thread exposed in Feng and Jin (2005) and combine it with Martin (2004). This introduction is in no way meant to convey the subtle art of a DFT practitioner, it is rather intended for the reader familiar with Anderson transition and general concepts of condensed matter physics.

2.1.1 Definition of the electronic structure problem

The central goal of condensed matter physics is to study the properties of solids. In particular, the conduction properties of a material can be understood by studying the permitted energy states of an electron inside it.

A solid is composed of many atoms,¹ each contributing to the total energy of the system with the kinetic energy of its nucleus and surrounding electrons, as well as with the Coulomb interaction between all charged subsystems. Assuming the absence of electromagnetic fields and ignoring the spin degrees of freedom,² we may write the Hamiltonian of this system as

$$H = \sum_i \frac{p_i^2}{2m} + \sum_\alpha \frac{P_\alpha^2}{2M_\alpha} + \sum_{i>j} \frac{e^2}{\|\mathbf{r}_i - \mathbf{r}_j\|} + \sum_{\alpha>\beta} \frac{Z_\alpha Z_\beta e^2}{4\pi\epsilon_0 \|\mathbf{R}_\alpha - \mathbf{R}_\beta\|} - \sum_{i,\alpha} \frac{Z_\alpha e^2}{4\pi\epsilon_0 \|\mathbf{r}_i - \mathbf{R}_\alpha\|},$$

where \mathbf{p}_i , \mathbf{r}_i , m and $-e$ indicate, respectively, the momenta, coordinates, mass and charge of the electrons, while \mathbf{P}_α , \mathbf{R}_α , M_α and $Z_\alpha e$ indicate the same quantities for the nuclei, and ϵ_0 is the dielectric constant of the vacuum. In the framework of quantum mechanics all positions and momenta, and hence the Hamiltonian, are operators in the Hilbert space describing the corresponding electron or nucleus. It is customary to cast the Hamiltonian operator in the position representation by substituting $\mathbf{p}_i \rightarrow -i\hbar\nabla_i$ and $\mathbf{P}_\alpha \rightarrow -i\hbar\nabla_\alpha$, and write the corresponding eigenvalue equation, i.e. the Schrödinger equation of the many-body interacting system that is our solid. This equation, enclosing all physical properties of the system, is practically impossible to solve, since the dimension of the solution space

¹Remember Avogadro's number $N_A \approx 6.022 \times 10^{23}$ atoms.

²Unless we need to compute the magnetic properties of a material.

grows exponentially with the number of atoms in the material. One of the goals of the theory of condensed matter is to find approximations that can be used to simplify this problem.

One approximation consists in regrouping the system in ions, composed of the nuclei and inner-shell electrons, and valence electrons, which can be delocalised and thus contribute to electric or thermal transport. This *valence electron approximation* (Hellmann, 1935) holds for many materials, e.g. alkalis or noble metals (Hellmann and Kassatotschkin, 1936), but does not apply when the valence electrons are in the inner shells, like in transition metals (*d* electrons) or rare earths (*f* electrons).

We then notice that the mass of the nuclei is, usually, several orders of magnitude larger than that of the electrons, which means that the time scale of the nuclear motion is also much longer. The electrons, moving correspondingly much faster, can thus *adiabatically* follow the movement of the nuclei. This separation of time scales allows³ to decouple the electronic from the ionic motion, known as *Born-Oppenheimer approximation* (Born and Oppenheimer, 1927). We can therefore write a separate Schrödinger equation for the electrons, which, in atomic units ($\hbar = m = e = 4\pi\epsilon_0 = 1$), reads

$$\left[-\frac{1}{2} \sum_i \nabla_i^2 + \sum_{i>j} \frac{1}{\|\mathbf{r}_i - \mathbf{r}_j\|} + \sum_i v(\mathbf{r}_i) \right] \Psi(\mathbf{r}_1, \dots, \mathbf{r}_N) = E\Psi(\mathbf{r}_1, \dots, \mathbf{r}_N). \quad (2.1)$$

In this approximation, all electrons move in the static potential v due to the ions, as happens in a periodic or homogeneous structure. The *electronic structure problem*, then, is the solution of the many-body Schrödinger equation (2.1) for the electronic degrees of freedom.

2.1.2 Solution via density functionals

While most textbooks, e.g. Feng and Jin (2005), would now follow the historical progression of introducing the Hartree and Hartree-Fock approximations, we will skip directly to *density functional theory* (DFT). The difference between these methods is revealed by the name: the starting point for the former is the wave function, while for the latter it is the *electron density*.

³It can be shown that the error on this approximation depends on $(m/M)^{1/4}$, where M is the mass of an ion.

The foundation of DFT relies on the theorems of Hohenberg and Kohn (Hohenberg and Kohn, 1964), which state that (i) the ground-state density $n_0(\mathbf{r})$ of any system of electrons uniquely determines the external potential $v_{\text{ext}}(\mathbf{r})$ acting on them, and that (ii) for any such $v_{\text{ext}}(\mathbf{r})$ there exists a functional E_{HK} of the electron density $n(\mathbf{r})$ that is minimised by the ground-state density $n_0(\mathbf{r})$ and whose minimum is the ground-state energy of the system. It goes beyond the scope of this work to prove these theorems, interested readers may consult any modern textbook on condensed matter (Feng and Jin, 2005; Patterson and Bailey, 2010), electronic structure (Martin, 2004), or DFT (Engel and Dreizler, 2011).

As is usually noted, the Hohenberg-Kohn theorems prove the existence of such a functional E_{HK} but do not give instructions as to its calculation. Nevertheless, there are further approximations we can make to simplify (2.1). We can split E_{HK} itself in three terms that follow those in (2.1), namely a kinetic term $T[n]$, a term for the interaction between the electrons $E_{\text{int}}[n]$, and a term $E_{\text{ext}}[n]$ for the interaction with an external field (e.g. that of the ions): $E_{\text{HK}}[n] = T[n] + E_{\text{int}}[n] + E_{\text{ext}}[n]$. Notice that this functional, evaluated at the ground-state density $n_0(\mathbf{r})$, yields the total energy of the electrons: if we want the total energy of the system we need to add a contribution from the ions, which will be a constant if they have fixed positions.

While we can simply define⁴

$$E_{\text{ext}}[n] = \int V_{\text{ext}}(\mathbf{r})n(\mathbf{r}) \, d\mathbf{r} ,$$

we need to use a further ansatz to get simpler expressions for E_{int} and T . In this framework, called *Kohn-Sham* DFT, we put aside our original system and take a set of non-interacting electrons with the *same* density $n(\mathbf{r})$. Logically, we now need to define a new functional E_{KS} that yields the same energy as that of the original system described by E_{HK} . As we are treating a system of indistinguishable non-interacting electrons, we can describe each of them with a local single-particle orbital $\psi_i(\mathbf{r})$ and construct the many-body state as a Slater determinant (Kohn and Sham, 1965). The corresponding density, therefore, has the simple form

$$n(\mathbf{r}) = \sum_i |\psi_i(\mathbf{r})|^2 , \quad (2.2)$$

⁴Notice that now the definition of the external potential is $V_{\text{ext}}[n] = \frac{\delta E_{\text{ext}}}{\delta n}$, as opposed to that in (2.1).

with the summation running over all electrons (and spins, if relevant to the problem). Inspired again by (2.1), we define two new terms: the kinetic energy of the new, non-interacting system

$$T_s[n] = -\frac{1}{2} \sum_i \int \psi_i^*(\mathbf{r}) \cdot \nabla^2 \psi_i(\mathbf{r}) d\mathbf{r}, \quad (2.3)$$

and the Coulomb interaction of the density of electrons with itself, which we call the *Hartree energy*

$$E_H[n] = \frac{1}{2} \iint \frac{n(\mathbf{r})n(\mathbf{r}')}{\|\mathbf{r} - \mathbf{r}'\|} d\mathbf{r} d\mathbf{r}'. \quad (2.4)$$

Since we have no reason to believe that the kinetic and interaction terms of the initial interacting “Hohenberg-Kohn system” and its equivalent non-interacting “Kohn-Sham system” add up to the same value, we need to introduce a term to account for the possible difference, the *exchange-correlation* (XC) functional $E_{xc}[n]$. We can now write the final form of the Kohn-Sham energy functional as

$$E_{KS}[n] = T_s[n] + E_H[n] + E_{xc}[n] + E_{ext}[n]. \quad (2.5)$$

2.1.3 The Kohn-Sham equations

With the Kohn-Sham approximation we have simplified the problem by introducing terms that can be computed exactly (T_s and E_H) and hiding everything else in the exchange-correlation term E_{xc} . There may be cases, e.g. for weakly-interacting electrons, where the magnitude of E_{xc} is much smaller than the other terms and an approximated form is sufficient. In fact, fundamental research in DFT deals with the formulation of appropriate functionals to solve the electronic structure problem for different materials (Neugebauer and Hickel, 2013).

We will not comment further on the role of E_{xc} , and instead simply continue on the path indicated by the Hohenberg-Kohn theorems. The ground state and its energy are found by minimising E_{KS} with respect to the single-particle orbitals ψ_i . We will simplify the calculation by differentiating with respect to the conjugate ψ_i^* , using the chain rule, and Eqs. (2.2) and (2.3):

$$\begin{aligned} \frac{\delta E_{KS}}{\delta \psi_i^*} &= \frac{\delta T_s}{\delta \psi_i^*} + \left[\frac{\delta E_H}{\delta n} + \frac{\delta E_{xc}}{\delta n} + \frac{\delta E_{ext}}{\delta n} \right] \frac{\delta n}{\delta \psi_i^*} \\ &= -\frac{1}{2} \nabla^2 \psi_i(\mathbf{r}) + [V_H(\mathbf{r}) + V_{xc}(\mathbf{r}) + V_{ext}(\mathbf{r})] \psi_i(\mathbf{r}). \end{aligned} \quad (2.6)$$

The potentials V_H , V_{xc} and V_{ext} are all defined as the functional derivative with respect to the density $n(\mathbf{r})$ of the corresponding energy functionals. Since E_{KS} is minimised under the constraint that the single-particle orbitals are normalised, we need to introduce a set of Lagrange multipliers ε_i , obtaining

$$\left[-\frac{1}{2}\nabla^2 + V_H(\mathbf{r}) + V_{xc}(\mathbf{r}) + V_{ext}(\mathbf{r}) \right] \psi_i(\mathbf{r}) = \varepsilon_i \psi_i(\mathbf{r}). \quad (2.7)$$

This is the *Kohn-Sham equation* and the operator in square brackets is the *Kohn-Sham Hamiltonian*. As a Schrödinger-like equation, Eq. (2.7) is solved as an eigenvalue problem. Notice that the effective potential $V_{eff}(\mathbf{r}) = V_H(\mathbf{r}) + V_{xc}(\mathbf{r}) + V_{ext}(\mathbf{r})$ still depends on the electron density, which is itself defined by the single-particle orbitals that solve (2.7). Like in the Hartree and Hartree-Fock methods, the Kohn-Sham equation must be solved self-consistently: we start from a trial density and compute V_{eff} , from which Eq. (2.7) is solved to find new single-particle orbitals ψ_i , from which a new density is calculated. This cycle is repeated until we find the initial and final densities to be equal within a certain threshold. Once convergence has been reached, those densities are equal to the ground-state density n_0 and we can obtain the corresponding energy $E_0 = E_{KS}[n_0]$.

We shall remark here that, while Eq. (2.7) is a Schrödinger-like equation, the single-state orbitals ψ_i are *not* the wave functions of the electrons and the ε_i 's are *not* the corresponding energies. The only exception is the energy of the highest occupied single-state orbital, which can be proved to be equal to the ionization potential of the interacting system (Janak's theorem—see Engel and Dreizler (2011) for proof). The reason is that the Kohn-Sham system we have built is an approximation of the original interacting problem that is meant to give the same electron density and *total* energy. The difference between the Kohn-Sham eigenvalues, however, can still be interpreted as a zeroth order approximation to the excitation energies of the system, as shown in Görling (1996). For this reason, we can compare the band structures calculated by DFT with that obtained in angle-resolved photoemission spectroscopy (ARPES) experiments, and thus assess whether the Kohn-Sham eigenvalues are a good approximation to the real energy bands (Giustino, 2014).

Other important repercussions of our approximations are that DFT strictly computes properties of the *ground state* only, and does not give any information on excited states without extending its original foundations (with e.g. the *GW* approximation). Furthermore, it only treats *uncorrelated* electrons, those whose

positions do not influence each other via the Coulomb repulsion. This does not contradict our premises: DFT still models the Coulomb interaction within V_{eff} , and it does indeed work for systems where electrons effectively only interact via a mean field. However, like all other independent-particle methods, traditional semi-local functionals fail to capture the physics of “strongly correlated” systems like Mott insulators or heavy-electron metals. Please refer to Martin (2004) for further discussion and references.

2.1.4 Pseudopotentials

In addition to the exchange–correlation functional, *ab initio* calculations using common basis sets, such as plane waves, rely also on the choice of a *pseudopotential*, namely a description of the Coulomb interaction of the nuclei. The potential around a nucleus is a central diverging potential that will scatter any particle travelling towards it. As shown in any quantum mechanics textbook, e.g. Cohen-Tannoudji et al. (2005), this problem can be solved by using a decomposition in partial waves. In the asymptotic regime, far away from the divergence, the particle behaves like a superposition of plane waves with a phase shift determined by the potential. We can therefore substitute the original nuclear potential with another potential that induces the same phase shifts and hence yields the same results away from the cores.

In the original derivation, presented e.g. in Feng and Jin (2005) or Patterson and Bailey (2010), the wave function for the valence electron is decomposed in a “pseudofunction” and a linear combination of atomic orbitals for the electrons in the inner shells. It is then shown that the pseudopotential that gives rise to the “pseudofunctions” via the Schrödinger equation is much weaker than the original potential. The divergences of the Coulomb interaction of the nuclei is compensated by the atomic orbitals of the core electrons, known as the *cancellation theorem* (Cohen and Heine, 1961). The weakness of the pseudopotential also proves the free-particle or delocalised behaviour of the valence electrons, except for the already-mentioned cases of transition metals and rare earths where there is no clear distinction with the core electrons.

2.1.5 Linear scaling DFT

One of the aspects that makes DFT so successful is its computational cost, which scales *polynomially* as $O(N^3)$ with the number of atoms N . This scaling compares

favourably to the original problem (2.1), whose solution space scales *exponentially* as 3^N , and also to other methods based on correlated wave functions, which usually scale with N^5 (Møller and Plesset, 1934) to N^7 (Coester and Kümmel, 1960). Despite this, a cubic scaling still makes it difficult to study systems with more than a few hundred atoms (e.g. a large molecule) even on national-level high-performance computing facilities. This limitation has prompted the research into *linear-scaling* DFT, that is methods to achieve $O(N)$ scaling of the computation. These methods exploit the so-called “near-sightedness” of quantum mechanics, namely the assumption that a part of the system is unaffected by a perturbation far away (Kohn, 1996). In practice, it is up to the user to define how far is actually “far”, where longer radii of influence yield more accurate, but also more computationally expensive calculations. We will now discuss how this is achieved in the specific implementation that we use throughout this project, called ONETEP (Haynes et al., 2006).

We start by extending the notion of electron density to a matrix $\hat{\rho}$, which, in the position representation, reads

$$\rho(\mathbf{r}, \mathbf{r}') = \sum_i f_i \psi_i^*(\mathbf{r}) \psi_i(\mathbf{r}'), \quad (2.8)$$

where $f_i \in [0, 1]$ is the occupation number of the Kohn-Sham orbital ψ_i . The electron density (2.2) is the diagonal of this density matrix: $n(\mathbf{r}) = 2\rho(\mathbf{r}, \mathbf{r})$, with the factor of 2 accounting for the spin degeneracy. In this formalism, the total energy of the system is $E = 2 \text{Tr}[\hat{\rho}\hat{H}]$, where \hat{H} is the Kohn-Sham Hamiltonian in the basis of $\{\psi_i\}$. The minimisation of E , under the constraints $N = 2 \text{Tr}[\hat{\rho}]$ (normalisation) and $\hat{\rho}^2 = \hat{\rho}$ (idempotency of a projection operator), is efficiently implemented following the method by Li et al. (1993). It has been shown (see Refs. in Haynes et al. (2006)) that for a material with a finite band gap (an insulator), the density matrix decays *exponentially* with the distance $\|\mathbf{r} - \mathbf{r}'\|$. This justifies the near-sightedness assumption and allows to truncate the density matrix elements for sites further apart than a cutoff r_K , such that the information contained in $\hat{\rho}$ scales linearly with the number of atoms.

Another feature that allows to achieve order- N scaling is the introduction of a set of support functions ϕ_α centred on each atom α , called *nonorthogonal generalized Wannier functions* (NGWFs). These functions, similar to local orbitals, extend only up to cutoff radii r_{NGWF} . This cutoff defines the “near-sightedness” of the DFT, since the local orbitals on two sites will overlap only if their distance is less than

$2r_{\text{NGWF}}$. Publications about ONETEP employ standard bra-ket notation and write $|\phi_\alpha\rangle$, where $\phi_\alpha(\mathbf{r}) = \langle \mathbf{r} | \phi_\alpha \rangle$. These support functions have corresponding *dual* functions ϕ^α such that

$$\langle \phi_\alpha | \phi_\beta \rangle = \int \phi_\alpha^*(\mathbf{r}) \phi_\beta(\mathbf{r}) d\mathbf{r} = S_{\alpha\beta} \quad \langle \phi_\alpha | \phi^\beta \rangle = \int \phi_\alpha^*(\mathbf{r}) \phi^\beta(\mathbf{r}) d\mathbf{r} = \delta_{\alpha,\beta} . \quad (2.9)$$

Here $S_{\alpha\beta}$ is the *overlap matrix* and represents the *metric tensor* of the space defined by a nonorthogonal basis, while $\delta_{\alpha,\beta}$ is the Kronecker delta. Notice that for an orthogonal basis $S_{\alpha\beta} = \delta_{\alpha,\beta}$ and we would identify the support functions with their duals. For further discussion please refer to O'Regan (2012). This implies, then, that every problem we solve with ONETEP is described not only by a Kohn-Sham Hamiltonian $H_{\alpha\beta}$, but also by an overlap matrix $S_{\alpha\beta}$ that describes the basis of NGWFs.

2.1.6 Forces and geometry relaxation

When we simulate a compound material, as in our case, the atom of the impurity species might deform the surrounding lattice. For this reason, the final topic we touch is how to relax the crystal structure to minimise the total energy of a system.

This process is also a cycle that starts with the self-consistent minimisation of the total energy when the ions are kept fixed. The positions of the atoms are then perturbed to find a lower total energy. We repeat the self-consistent cycle to find the ground state and then compute the force on the i -th ion via the Hellmann-Feynman theorem as

$$\mathbf{F}_i = -\frac{dE}{d\mathbf{r}_i} = -\langle \Psi | \frac{dH}{d\mathbf{r}_i} | \Psi \rangle . \quad (2.10)$$

We finally update the positions and start again until the changes in energy and forces are within a specified threshold. Notice that we solve the Kohn-Sham equations self-consistently every time we change the positions of the ions, although the converged electronic structure information of one step is used as starting guess in the next. This ensures that the electronic and ionic degrees of freedom are always decoupled, as required in the Born-Oppenheimer approximation, but also implies that a geometry relaxation requires significantly more resources than a fixed-geometry calculation. Its practical implementation is possible using methods like the Broyden-Fletcher-Goldfarb-Shanno algorithm (Press et al., 2007), which

relies on an iteratively-improved approximation of the inverse Hessian matrix.

2.2 Fundamentals of Si: crystal structure, semiconducting properties, simulations

In this section we point out the necessary details to simulate pure or doped silicon with ONETEP.

2.2.1 Diamond cubic structure

At standard pressure, silicon crystallises in a diamond cubic lattice, like carbon and other group-IV elements. The tetrahedral arrangement of the bonds to its four nearest neighbours can be explained by the formation of hybridized sp_3 orbitals, where one electron is promoted from the $3s$ to the $3p$ shell (Yu and Cardona, 2010). As a result, silicon atoms have four valence electrons, instead of two. A unit cell of the diamond cubic lattice contains 8 atoms (see Fig. 2.2, right), whose positions are given by

$$\left\{ \frac{a}{4}(x, y, z) : x, y, z \in \mathbb{Z}, (x \equiv y \equiv z \pmod{2}) \wedge (x + y + z \pmod{4}) \in \{0, 1\} \right\}, \quad (2.11)$$

where $a \approx 5.4 \text{ \AA}$ is the lattice constant (Nagy and Strand, 2009). The explicit coordinates of the points, in units of $a/4$, are

$$1 : (0, 0, 0), (2, 2, 0), (2, 0, 2), (0, 2, 2) \quad 2 : (3, 3, 3), (1, 1, 3), (1, 3, 1), (3, 1, 1), \quad (2.12)$$

which can be seen as two interpenetrating face-centred cubic (fcc) lattices, which we label '1' and '2'. Notice that the two sublattices are separated by $a/4(1, 1, 1)$, hence no translation $a(m, n, p)$, with $m, n, p \in \mathbb{Z}$, can transform a point of 1 (resp. 2) into one of 2 (resp. 1). We use this translation vector to build the *super cell*, i.e. the lattice of a larger system composed of several unit cells in the three dimensions.

2.2.2 Simulations of pure Si

In Fig. 2.1 we show the structure and essential keywords of a ONETEP input file used to simulate a system of 4096 Si atoms. For the meaning of the single keywords we refer the reader to the documentation on the ONETEP website,⁵ and to Skylaris

⁵<http://www.onetep.org/Main/Keywords>

```

task                : SINGLEPOINT
xc_functional       : PBE
kernel_cutoff       : 1000.0 bohr
write_hamiltonian   : T
maxit_ngwf_cg       : 18
ngwf_threshold_orig : 1.4e-6
maxit_lnv           : 30
minit_lnv           : 10
lnv_threshold_orig  : 2e-7
elec_energy_tol     : 1e-7 hartree
cutoff_energy       : 880.0 eV
psinc_spacing       .5083350000 .5083350000 .5083350000

%block species
Si Si 14 9 8.0
%endblock species

%block species_atomic_set
Si "SOLVE ws=2 wp=2.2 wd=6.6 S=400 R=8"
%endblock species_atomic_set

%block species_pot
Si "silicon.recpot"
%endblock species_pot

%block lattice_cart
81.3336  0.0000  0.0000
  0.0000 81.3336  0.0000
  0.0000  0.0000 81.3336
%endblock lattice_cart

%block positions_abs
Si      0      0      0
Si      5.08335      5.08335      0
...
%endblock positions_abs

```

Figure 2.1: Essential blocks and keywords of the ONETEP input file to simulate pure Si.

and Haynes (2007) for a discussion on the choice of the different parameters and thresholds for crystalline Si.

The important point that needs to be mentioned in this context concerns the choice of NGWFs. Silicon is an element of the IV group, has 14 electrons, and its electronic configuration reads $[\text{Ne}]3s^23p^2$. The minimal basis set, then,

would consist of the single s and the three p orbitals, for a total of four NGWFs. Convergence with this basis set, however, requires rather large truncation radii (usually $r_{\text{NGWF}} = 12a_0$, where $a_0 = \hbar^2/(me^2)$ is the Bohr radius). The solution we instead adopt to achieve better convergence is to expand the basis set with the five d orbitals, and reduce their truncation radii to $r_{\text{NGWF}} = 8a_0$. All the matrices throughout this project will be written in this basis, hence all of them will have a linear size of $9N$, if N is the number of atoms.

The other point that is worth mentioning is that the kernel cutoff we have chosen is $r_K = 1000a_0$, which is much larger than the box size $81.1336a_0$. This means that we are not imposing any cutoff beyond the near-sighted description arising from the NGWFs (see Sec. 2.1.5).

In Fig. 2.2 we show the probability amplitude $|\psi(\mathbf{r})|^2$ of the electron in the highest occupied Kohn-Sham eigenstate for a system of 4096 Si atoms. This state is at the top of the valence band and, as expected from a periodic potential, it is extended like a Bloch wave throughout the volume.

By construction, the band gap calculated in Kohn-Sham DFT cannot equal the true band gap (Martin, 2004). In practice DFT gives very good qualitative predictions, but quantitatively underestimates the band gap, one reason being that it underestimates the correlations between electrons (Feng and Jin, 2005). The estimate obtained in our simulation is 0.691 eV (Fig. 2.6), which is 59% of the experimental value of 1.17 eV at 4 K (Green, 1990).

2.2.3 Simulation of doped Si

As discussed before, we can subdivide the diamond lattice into ‘1’ and ‘2’ sublattices according to Eq. (2.12). If we imagine standing at a site from ‘1’, like $(0, 0, 0)$, we will see our nearest neighbours at the displacements⁶

$$\Delta_1 = \{(-1, -1, -1), (-1, 1, 1), (1, -1, 1), (1, 1, -1)\}. \quad (2.13)$$

On the other hand, if we are sitting at a ‘2’ site, our neighbours will be at

$$\bar{\Delta}_1 = \{(1, 1, 1), (1, -1, -1), (-1, 1, -1), (-1, -1, 1)\}, \quad (2.14)$$

which is the point reflection of Δ_1 . Because a reflection in three dimensions is not orientation-preserving, namely it cannot be replicated by a simple rotation, we

⁶In the rest of this section we will be writing lengths in units of $a/4$.

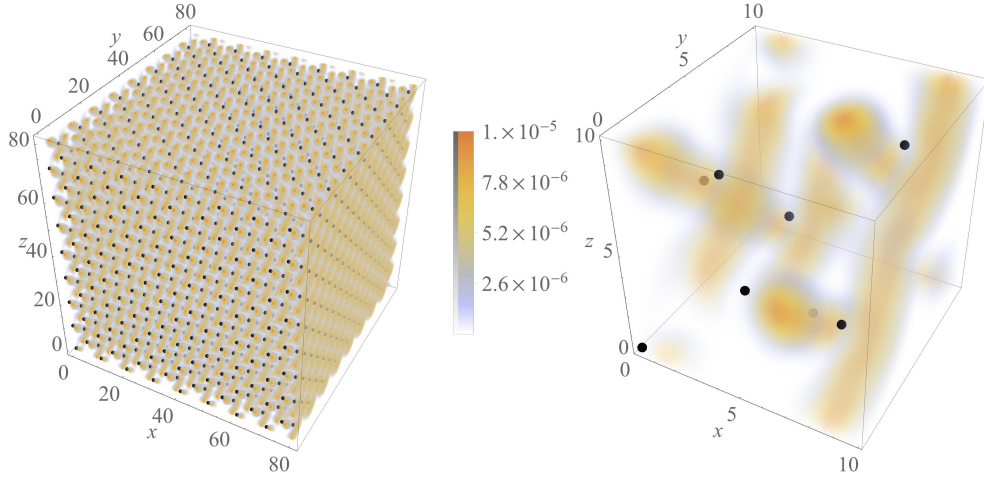


Figure 2.2: Probability amplitude $|\psi(\mathbf{r})|^2$ of the highest occupied Kohn-Sham eigenstate of 4096 Si atoms, as obtained from our DFT simulation. On the left (right) we show the full system (a unit cell). Small (high) amplitudes are indicated in opacity and colour by transparent violet (solid red). Black circles indicate the lattice positions. Lengths are expressed in Bohr radii.

need to simulate an isolated impurity both in ‘1’ and ‘2’ positions. The reason is that dopants in the different sublattices form bonds with Si in opposite directions, and the corresponding matrix elements, written in the basis of atomic-orbital-like support functions, would need to be rotated accordingly, which is non-trivial.

This observation also determines how many cases we need to simulate if we want to include pairs of impurities. If we consider up to the fourth shell of neighbours, as we justify in Sec. 3.2 and Fig. 3.3 for this project, an atom has 34 neighbours. This would lead to a total of 68 simulations (34 configurations for each sublattice), but we now show that only half of them are needed. Odd shells connect atoms between ‘1’ and ‘2’, hence we only need one simulation for each displacement in (2.13), for instance that with the impurities at $(0, 0, 0)$ and $(-1, -1, -1)$. Even shells, instead, connect sites from the same sublattice, since the relative orientation of the orbitals might be different. Because of the same symmetry, however, to each displacement δ there corresponds a $-\delta$, so we only need to do half of the simulation for each sublattice. In Fig. 2.3 we have reported the displacement vectors for the first two shells (for an atom in the ‘1’ sublattice) as an example. For the second shell, displacements in opposite directions have been written in adjacent columns. In total, then, we need to do 34 simulations for the pair configuration,

SHELL I								
-1/4	-1/4	-1/4	-1/4	1/4	1/4	1/4	1/4	-1/4
SHELL II								
-1/2	-1/2	0	1/2	1/2	0	-1/2	1/2	0
0	-1/2	-1/2	0	1/2	1/2	0	1/2	-1/2
-1/2	0	1/2	1/2	0	-1/2	-1/2	0	-1/2

Figure 2.3: Displacement vectors connecting a ‘1’ site to its first and second nearest neighbours in the diamond cubic lattice, expressed in units of lattice parameters.

in addition to the 2 simulations for the single impurity case.

Compared to Fig. 2.1, the ONETEP input file for the doped samples should consistently use the same thresholds and cutoffs. The only needed change is in the species, species_atomic_set and species_pot blocks, where the corresponding lines for the dopant species should be added. For all simulations that have been used as basis material for Ch. 3, we have tightened the convergence threshold `lnv_threshold_orig` to 10^{-9} to obtain more accurate occupancies (at least 0.99%) of the defect states. Each fixed-geometry simulation was run on 1152 cores for around 12 h on ARCHER, the UK National High Performance Computing Facility.

As we will see in the following section, if we consider an impurity that induces a lattice deformation, we also need to geometrically relax the lattice for each configuration. In this case, we can save computational resources by calculating the distortion for one displacement and reusing it as a starting configuration for the other displacements to the same shell. For the same symmetry arguments derived from having two distinct sublattices, one may reuse a distorted lattice by rotating and/or inverting it with respect to the centre of the simulation box. The reader who attempts this should make sure that the undistorted lattice is correctly mapped onto itself! A geometry relaxation simulation, also on 1152 cores, can take up to 5 times more using a tolerance on the maximum atomic forces of 1 eV/\AA . As we show in Sec. 2.3, the relaxation of the lattice leads to a decrease in the energy of the impurity states appearing in the band gap. We have not considered a tighter constraint on the atomic forces, as this would require a significant increase in computational costs that would not lead to an appreciable change in the properties of the doped system.

2.3 Doping, shallow impurities and deep centres

The conduction properties of a semiconductor can be modified by the introduction of *defects*. While there are several types of defects, we focus now on *substitutional* defects, which occur when a host atom is replaced by an atom of a different species. When such defects contribute free electrons to the material, they are called *donors*, when they remove them (or “contribute holes”) they are called *acceptors*.

For silicon, an element from group IV, donors are elements from group V, i.e. P, As or Sb, while double donors are elements from group VI, chalcogens like S, Se or Te. Examples of acceptors, instead, include group-III elements like B, Al, Ga or In, while group II provides double acceptors like Be or Zn. In the following subsections we review the effects of doping Si with P and S, the elements we have considered for this project (see Sec. 3.5).

2.3.1 Shallow impurities: doping with phosphorus

Let us consider a system of pure Si and then substitute one Si for a P atom. In terms of charges, the new atom adds one valence electron and one proton in the nucleus. This additional electron, which is not used to form covalent bonds with the neighbouring atoms, is free to move in the ionic potential of P, but is screened by the presence of the other Si atoms. A good approximation for the resulting effective potential, then, is that of a hydrogen atom in an effective medium, silicon (Resta, 1986; Feng and Jin, 2005; Yu and Cardona, 2010). Therefore, the permitted energies for the donated electron are, like for the hydrogen atom, described by the Rydberg series:

$$E_n = E_{CB} - \frac{R_*}{n^2} \quad R_* = \frac{m_*}{m} \frac{1}{\epsilon^2} R. \quad (2.15)$$

Here R_* is the Rydberg constant for the donated electron, which is proportional to the Rydberg constant for the hydrogen atom $R = e^4 m (2\hbar^2)$ via the effective mass m_* and the dielectric constant ϵ of the medium. The ground state has index $n = 1$ and the energies of the excited states are closer to the conduction band energy E_{CB} with higher n . For this reason, donor impurities are called *hydrogenic* or *shallow*.

The ground state itself is a localised state whose envelope function can also be

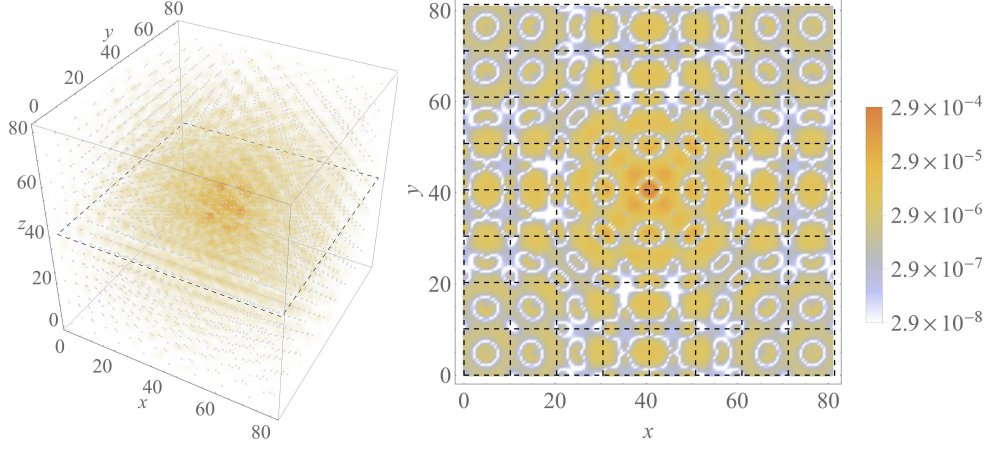


Figure 2.4: Probability amplitude $|\psi(\mathbf{r})|^2$ of the impurity state induced by a P impurity in the 4096 Si system, as obtained from our DFT simulation. On the left we show the whole cell, while on the right we show a cut at $z \approx 40.67a_0$ (indicated with a dashed line on the leftward picture). The logarithmic scale on the right holds for both figures, and in the leftwards plot we also scale opacity from transparent (small values) to solid (higher values). The 10% smallest values are not shown. For the full system we show the lattice positions in light grey, while for the slice we separate the unit cells with dashed lines. Lengths are expressed in Bohr radii.

calculated (Feng and Jin, 2005; Yu and Cardona, 2010) as

$$F(r) = \frac{1}{(\pi a_*^3)^{1/2}} e^{-r/a_*} \quad a_* = \left(\frac{m}{m_*} \epsilon \right) a_0, \quad (2.16)$$

where, like for the Rydberg constant, we have introduced an effective Bohr radius a_* for the donated electron. For a free electronic charge in Si, the conduction-band effective mass is $m_* \approx 0.3m$, the scaled Rydberg constant is $R_* \approx 31$ meV, and the effective Bohr radius is $a_* \approx 38a_0$ Yu and Cardona (2010).

Following Sec. 2.2.3, we use ONETEP to simulate a system of 4096 Si atoms with one substitutional impurity. In Fig. 2.4 we plot the probability amplitude of the impurity state, singled out by its energy in the band gap. The corresponding energy gap with the conduction band of the bulk system is 16 meV (Fig. 2.6), which is in line with the underestimation of the band gap discussed in 2.2.2 and with the results by Smith et al. (2017).

2.3.2 Deep centres: doping with sulphur

Traditionally, *deep centres* were those defects that introduced energy levels near the middle of the band gap, although nowadays the notion has been extended to any defect that does not behave like a shallow impurity. Deep centres have more localised wave functions and introduce a highly-localised potential in the material, e.g. due to the difference in electronegativity between the impurity and the host material (Yu and Cardona, 2010).

Paralleling the discussion on P in Si, a sulphur impurity introduces two protons in the nucleus and two additional valence electrons, which it donates to the system. This system cannot be described with a hydrogen-like Hamiltonian and therefore, for deep centres, we cannot formulate an effective mass theory. Determining the potential in this context is non-trivial and beyond the scope of this thesis. In fact, the energy levels for deep centres are usually determined using first-principles approaches or the tight-binding method, which is well suited for the problem of an electron bound to a strong potential (Jaros, 1980; Bachelet, 1986; Yu and Cardona, 2010).

When calculating the energy levels introduced by deep centres, we need to be careful about the *lattice relaxation*: both the defect and the surrounding host atoms will move if the energy gained by distorting the lattice is less than the energy lost when the defect becomes a deep centre (Yu and Cardona, 2010). Otherwise, when the distortion is not energetically favourable, the defect will induce an energy level much closer to the conduction band, and will then be classified as a shallow impurity.

In Fig. 2.5 we show the probability amplitude of an electron in the impurity state induced by a deep centre, as calculated using ONETEP and allowing the lattice to relax. Compared to Fig. 2.4, the state is much more localised (notice the different scales in the legends). The energy to the conduction band of the bulk system is 0.150 eV (Fig. 2.6), which is almost 10 times bigger than for a P impurity. For comparison, the same energy for a fixed-geometry system is 0.129 eV: higher, as expected. In our simulations this energy level is doubly-occupied by the extra electrons donated by the S defect, as this is the configuration with the lowest energy (Overhof et al., 1991).

The extent to which the lattice is deformed is shown in Fig. 2.7. After around $15a_0$, the 9th shell of nearest neighbours, the distortion of the lattice looks negligible; in fact, beyond the first neighbours it is already an order of magnitude smaller.

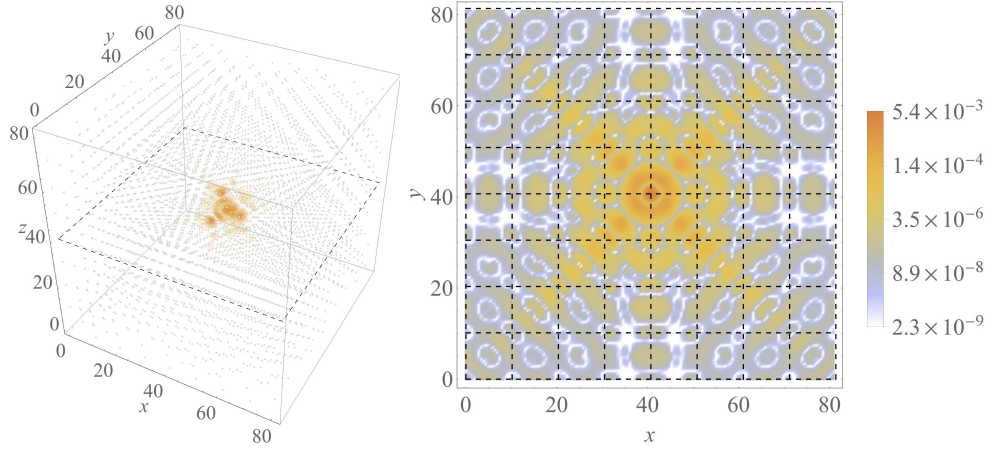


Figure 2.5: Same plots as in Fig. 2.4, but for a silicon system doped with one S impurity.

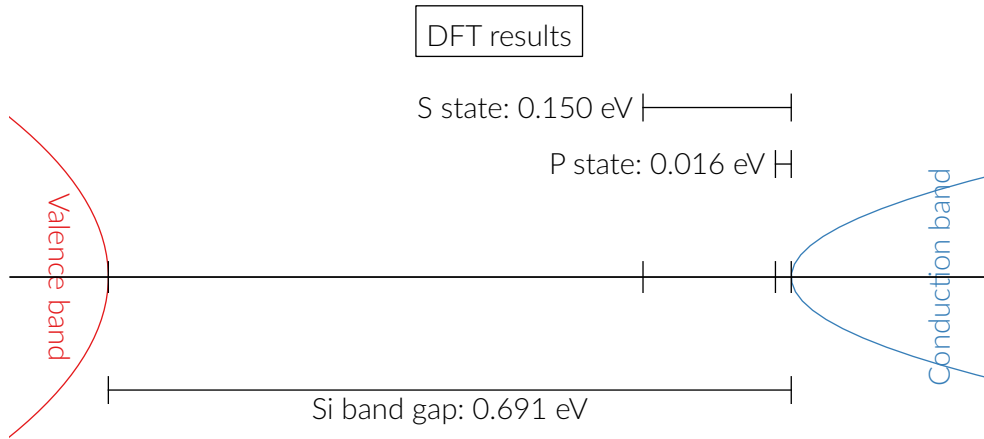


Figure 2.6: Summary of the results obtained with ONETEP on the band gap in Si and the impurity states of P and S. The parabolas represent the valence and conduction bands. The ticks indicate, from left to right, the top of the valence band, the S impurity state, the P impurity state and the bottom of the conduction band.

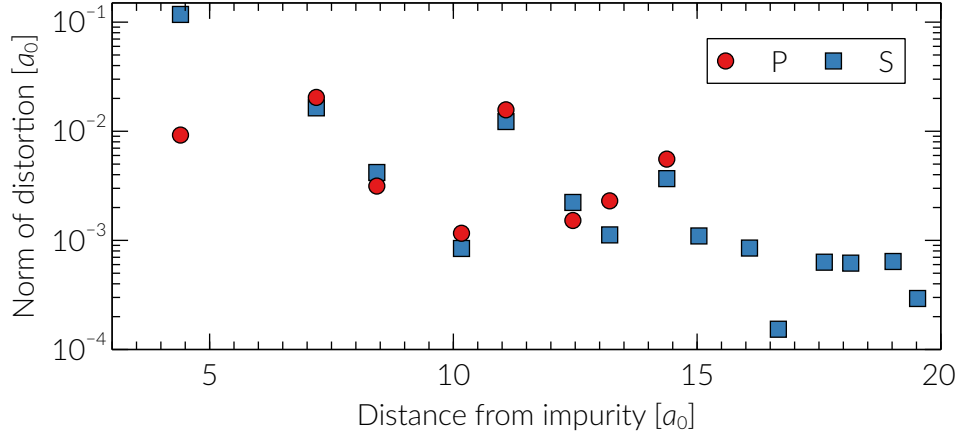


Figure 2.7: Magnitude of the distortion of the lattice introduced by a P (red circles) or S (blue squares) impurity in Si. We group the atoms, by their distance r to the impurity, in shells of nearest neighbours (abscissae). The atoms in each shell are displaced radially from the defect to a new distance r' . On the ordinates we show the norm of the displacement $\delta = |r - r'|$, which we find to be the same for all atoms in a shell (up to around 0.001%, due to the different rotational symmetry of the lattice and the local orbitals). For P, we find that the lattice is not distorted beyond the eighth shell of nearest neighbours.

2.4 Defect pair formation and doping techniques

In this section we discuss how the total energy of the system changes when we introduce two impurities at varying distance. The energy U needed to form a pair of defects at distance d is given by

$$U(d) = [E_{XX}(d) - E_{Si}] - 2(E_X - E_{Si}) = E_{XX}(d) - 2E_X + E_{Si}. \quad (2.17)$$

Here E_{Si} is the total energy to put together a system of pure Si, while E_X and E_{XX} , computed by ONETEP, are the total energy needed to build a system with, respectively, one and two impurities. The energy U , then, represents an “effective potential” between impurities that determines the distance separating two defects.

In Fig. 2.8 we show the results from P and S. For phosphorus, U is always positive, seems to oscillate but essentially vanishes for defects beyond the fourth shell of nearest neighbours (one unit cell). We can interpret this trend by saying that it is unlikely to find two P defects closer than a unit cell. For sulphur, instead, U is negative for direct neighbouring defects and then oscillates around zero for the other shells of neighbours. This implies that it is energetically favourable to

have two S impurities on adjacent sites, as we will further discuss in Sec. 3.2.

2.4.1 Growth of doped samples

The behaviour determined from U is related to the experimental doping techniques for Si. The metal-insulator transition in Si doped with P (Si:P) occurs at concentrations around 10^{18} cm^{-3} , two orders of magnitude lower than the corresponding solid solubility of roughly 10^{20} cm^{-3} (Bachelet, 1986). In very general lines, the growth of a doped semiconductor like Si:P involves the melting of the host material and the deposition on its surface of the impurities from a surrounding gas. During the resolidification process, the P atoms have enough time to diffuse and redistribute in a spatial configuration compatible with the potential defined by U .

For Si:S, instead, we have the opposite situation: the transition occurs at concentrations of 10^{20} cm^{-3} (Winkler et al., 2011), which is four orders of magnitude higher than the solid solubility of S in Si, around 10^{16} cm^{-3} (Bachelet, 1986). This implies that, if we used the same doping technique as with Si:P, the excess sulphur would precipitate. In this case, then, we need to *hyperdope* Si with S, as realised for the first time by Winkler et al. (2011) using ion implantation to shoot the defects into the host, and pulsed-laser melting and rapid resolidification to reform the surrounding lattice. With this technique, the behaviour deduced from U does not matter any more, as the S defects will be distributed randomly.

For more information on the growth of doped semiconductors please consult, e.g., Yu and Cardona (2010).

2.4.2 Generation of disordered samples

Crystal growth is a non-equilibrium process whose simulation is a non-trivial problem that we are not going to tackle or even discuss in depth here. The interested reader could start, e.g., from Landau and Binder (2000) and the references therein. To generate disordered samples of Si:P we run a simple Monte Carlo (MC) simulation using the Metropolis method: we generate an initial configuration with a fixed number of impurities and compute the total energy from the effective potential U . We then swap the position of an impurity with a neighbouring Si site and calculate the change in energy ΔE . Finally, we generate a random number $\alpha \in (0, 1)$ and, if $\alpha < e^{-\beta \Delta E}$, we accept the move. Here $\beta^{-1} = k_B T$, where k_B is Boltzmann's constant and T is the temperature at which the growth process




Figure 2.8: Energy needed to form a defect pair at a specific distance, for both P (red circles) and S (blue squares), obtained from ONETEP runs with 4096 atoms. For S we only show the results from the relaxed-geometry runs used in the later chapters (so up to the fourth nearest neighbours). The ordinates axis is broken in order to better show the difference in scale. Both data series are obtained while keeping the lattice fixed, which means that the points for S are an upper bound to the minimum total energy.

is happening. By repeating this procedure several times over all impurities in the system, we should reach a ground-state configuration.

As mentioned before, we assume that the hyper-doping technique used to grow S-doped samples is able to place those impurities randomly in the host. For this reason, we generate disordered realisations of this material by just changing the atom assigned to randomly picked sites from Si to S.

Chapter 3

From *ab initio* to effective models

In this Chapter we talk about the second part of the work flow, namely how we use the output from the *ab initio* simulations to construct effective models for many disorder realisations.

We start by recalling notions about the tight-binding model and connecting it to the Anderson model and the DFT description of the material. With these concepts, we can introduce our effective models, the underlying approximation and their technical implementation. We also highlight the aspects that our model retains from an atomistic description and set it apart from the paradigm of the Anderson model.

Finally, we discuss how to diagonalize the matrices to obtain eigenvalues and eigenvectors. We also point out that the several layers of complexity in our model make it a much harder problem than the Anderson model.

3.1 The tight-binding model

The tight-binding model (TBM) is used, as the name suggests, to describe electrons bound to a nucleus via a strong potential, strong enough that the interaction with the rest of the lattice can be considered a negligible perturbation (Diu et al., 1989). For this reason, the TBM has been widely used to study the electronic structure of deep centres, as mentioned in Sec. 2.3.2. The reader can find a discussion of the TBM in most textbooks on condensed-matter physics, including Kittel (2005), which we follow in the rest of the section.

An electron in said potential is described by an atomic orbital $\varphi(\mathbf{r} - \mathbf{r}_j)$ centred on the nucleus at \mathbf{r}_j , namely its wave function decays exponentially on length

scales much smaller than interatomic distances. The assumption behind the TBM, then, is that we can write the wave function of an electron *in a crystal* as a linear combination of said atomic orbitals:¹

$$\psi_{\mathbf{k}}(\mathbf{r}) = \frac{1}{\sqrt{N}} \sum_j e^{i\mathbf{k} \cdot \mathbf{r}_j} \varphi(\mathbf{r} - \mathbf{r}_j). \quad (3.1)$$

This is a Bloch function with wave vector \mathbf{k} for a crystal of N atoms, since we can show that $\psi_{\mathbf{k}}(\mathbf{r} + \mathbf{R}) = \exp(i\mathbf{k} \cdot \mathbf{R})\psi_{\mathbf{k}}(\mathbf{r})$, where \mathbf{R} is a translation from one lattice point to another.

Switching to bra-ket notation, we now define $\psi_{\mathbf{k}}(\mathbf{r}) = \langle \mathbf{r} | \psi_{\mathbf{k}} \rangle$ and $\varphi(\mathbf{r} - \mathbf{r}_j) = \langle \mathbf{r} | \varphi_j \rangle$. The expectation value of the Hamiltonian H is therefore

$$\langle H \rangle = N^{-1} \langle \psi_{\mathbf{k}} | H | \psi_{\mathbf{k}} \rangle = N^{-1} \sum_{j,m} e^{i\mathbf{k} \cdot (\mathbf{r}_m - \mathbf{r}_j)} \langle \varphi_j | H | \varphi_m \rangle. \quad (3.2)$$

Since we are working with a lattice, we can define $\boldsymbol{\rho}_j = \mathbf{r}_j - \mathbf{r}_m$ and simplify the double sum as

$$\langle H \rangle = \sum_j e^{-i\mathbf{k} \cdot \boldsymbol{\rho}_j} \int \varphi^*(\mathbf{r} - \boldsymbol{\rho}_j) H \varphi(\mathbf{r}) d\mathbf{r}. \quad (3.3)$$

At this point, it is customary to distinguish the diagonal elements ε (where $\boldsymbol{\rho}_j = 0$) and the off-diagonal $t(\boldsymbol{\rho}_j)$ as

$$\varepsilon = \int \varphi^*(\mathbf{r}) H \varphi(\mathbf{r}) d\mathbf{r} \quad \text{and} \quad t(\boldsymbol{\rho}_j) = \int \varphi^*(\mathbf{r} - \boldsymbol{\rho}_j) H \varphi(\mathbf{r}) d\mathbf{r}. \quad (3.4)$$

These terms of the TBM are also called, respectively, “self-energies” and “hopping” terms.

The Anderson model of Ch. 1 is the TBM of a simple-cubic crystal of hydrogen nuclei, i.e. of a material with only 1s atomic orbitals. In this case, however, we allow the diagonal elements ε_j to be different for each site, and we consider constant off-diagonal $H_{ij} = t$ if sites i, j are first nearest neighbours $\langle i, j \rangle$ and 0 otherwise. The corresponding Hamiltonian hence reads

$$H = \sum_j \varepsilon_j |\varphi_j\rangle \langle \varphi_j| + t \sum_{\langle i, j \rangle} |\varphi_j\rangle \langle \varphi_i|. \quad (3.5)$$

In the Anderson model, we usually set $t = 1$, while ε_j is a random variable

¹The TBM is also known, for this reason, as the Linear Combination of Atomic Orbitals (LCAO) approximation.

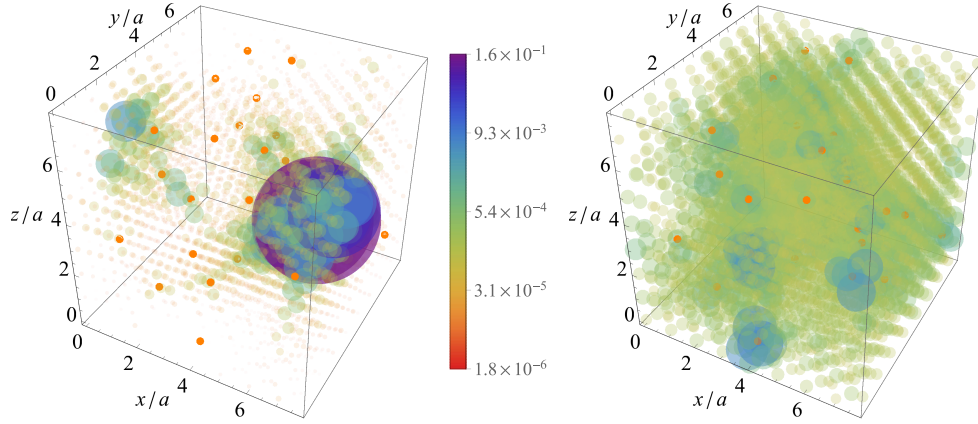


Figure 3.1: Wave functions for an exemplary systems of 4067 Si atoms and 29 S impurities. On the left we show a localised state deep in the impurity band and on the right an extended state above ε_F . We have represented the top 90% wave function values ψ_i with spheres of volume proportional to $|\psi_i|^2$. Opacity and colour are proportional to $-\log_L |\psi_i|^2$, with $L = 16$ here, so that lower (higher) values are in red transparent (violet solid). The box sizes are in units of a .

drawn from a uniform distribution over the interval $[-W/2; W/2]$, where the parameter W is the disorder strength. At a critical value $W_c \sim 16.5$ (Slevin and Ohtsuki, 1999) the wave function at the band centre $E = 0$ undergoes a transition from extended ($W < W_c$) to localised ($W > W_c$). While localisation can be defined more rigorously (del Rio et al., 1995), for this work we are happy with the commonly accepted form

$$|\psi(\mathbf{r})| \propto e^{-\frac{|\mathbf{r}-\mathbf{r}_0|}{\xi}}, \quad (3.6)$$

where \mathbf{r}_0 is a localisation centre and ξ is the localisation length introduced in Ch. 1. States that do not satisfy (3.6) are called interchangeably extended or delocalised.

Finally, we observe here that the Kohn-Sham Hamiltonian (2.7), written in the basis of NGWFs defined in (2.9), has the same structure as the TBM. The most important difference here is that each site is described by nine orbitals, instead of one. This implies that the Hamiltonian is divided in 9×9 -sized submatrices of self-energies (of and between the single NGWFs) and hopping terms (between NGWFs of the same or different sites). The overlap matrix has an analogous structure.

3.2 Translation from DFT to effective models

As mentioned in Ch. 1, simulating *ab initio* a sufficiently large sample of disorder realisations, for a number of concentrations and systems sizes, is simply not feasible, even with the efficient scheme implemented in ONETEP. Since they have the same structure, however, we can translate the DFT matrices into a TBM that is more efficient to build.

The idea of combining DFT and TBM is far from new, with early work including Andersen and Jepsen (1984), Sankey and Niklewski (1989) or Porezag et al. (1995). In those works the goal is to calculate *ab initio* specific tight-binding elements, an operation that usually requires the transformation from the standard-DFT basis of plane waves, to the localised basis set of the TBM.

In our approach, instead, we recognise the tight-binding structure of the Kohn-Sham Hamiltonian, when written in a basis of localised support functions, and we use it to build effective Hamiltonians that can reproduce the results obtained from the DFT. The focus is on building a DFT-like system rather than a TBM. This is achieved by simulating prototypes of pure or doped Si, as described in Ch. 2. The idea is the following: when we build the Hamiltonian² of a disordered sample, we assume that the self-energies and hopping blocks for each impurity are those calculated with ONETEP for a doped system. The final Hamiltonian, therefore, looks like the Hamiltonian for pure Si, except for those blocks that describe the presence of a defect. The matrix blocks that we reuse from the DFT simulations are stored in *catalogues*, whose construction and use we now describe (see Fig. 3.2).

The assumption behind using catalogues to build effective models is that the potential around each impurity is *locally* the same. It is up to us to decide the extent of this ‘locality’: the higher the range of the hopping, the more accurately we capture the effect of the impurity potential. For our effective models, we keep all hopping terms up to the 10th shell of nearest neighbours,³ since this is the extent of the “near-sighted” description via the local orbitals (see Sec. 2.1.5).

This picture should be completed by treating pairs of impurity with a separate catalogue, unless we can justify that (i) two defects do not interact at any distance or (ii) the formation of pairs is discouraged by an associated energy cost, as described in Sec. 2.4. For this reason, we have computed a catalogue of pairs only

²In the rest of the section, everything that is said about the Hamiltonian also holds for the overlap matrix.

³From here on, by “shell” we will always mean a shell of nearest neighbours.

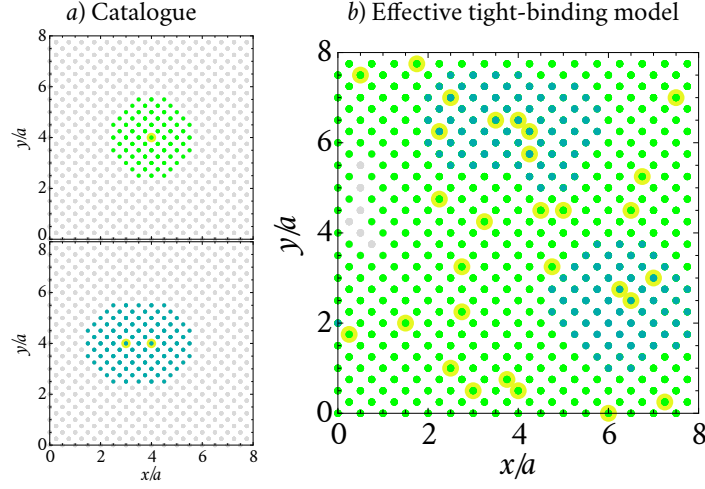


Figure 3.2: Construction of the effective tight-binding model. Part (a) represents the catalogue of prototypes. For clarity we show a projection on the xy plane and distances in units of a , the Si lattice parameter. The upper plot depicts one impurity (yellow) and the neighbouring Si atoms (green); the lower plot shows two impurities at distance a and their Si neighbours (dark green). Gray sites indicate Si atoms unaffected by the impurity potential. In (b) we show how we build an effective tight-binding model with 29 impurities. The colour code is the same as in (a) and indicates which catalogue is used. Due to the projection on the xy plane some impurities appear closer than they are. Reproduced from Carnio et al. (2017).

for sulphur, and not for phosphorus. Analogously to how we use the catalogue for isolated impurities, when a site is near a pair the Hamiltonian blocks are taken from the DFT simulation of that pair. For Si sites that are close to two impurities that are not considered a pair, we use the matrix blocks that connect it to the closest of the defects.

In this work, two defects are a pair when they are no more than a unit cell apart. This cut off can be decided by comparing the ONETEP runs of pairs at increasing distance to their effective model using the single-impurity catalogue. Because each defect induces a state in the band gap of Si, a parameter we can compare is the difference in energy between the defect states, see Fig. 3.3. When S defects are first nearest neighbours, they form bonding and anti-bonding states, where the former descend into the valence band while the latter remain in the band gap. In this case we compare the distance of the anti-bonding state to the lowest-energy conduction band state. While the 1-impurity catalogue manages to capture this feature of first nearest neighbours, the corresponding 2-impurity catalogues gives

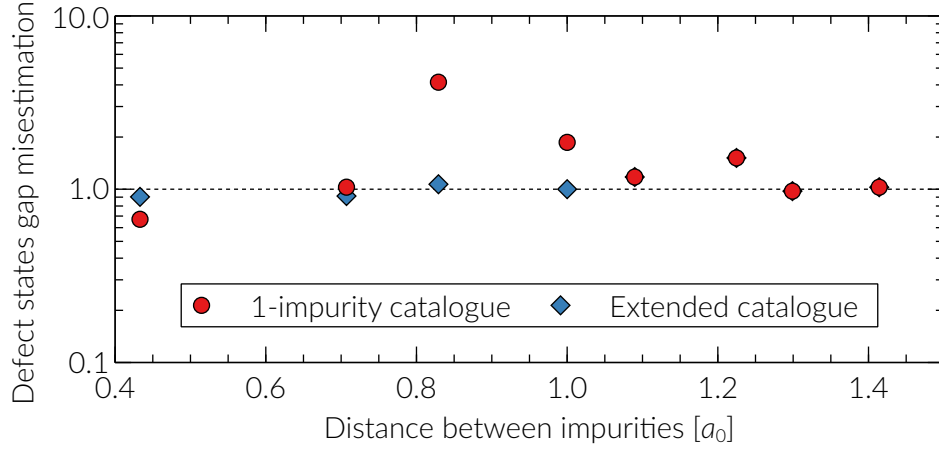


Figure 3.3: Ratio of the energy difference between the two impurity states appearing when a system of 4096 atoms is doped with a pair of defects at increasing distance. The ‘1-impurity catalogue’ is built from systems with only one impurity, while the ‘extended catalogue’ includes the description of pairs of defects. For situations where only one state in the gap appears, discussed in the text, the gap is taken between said state and the lowest conduction band state.

a better description by definition. Moreover, the 1-impurity catalogue predicts a similar situation when S defects are second nearest neighbours. This contradicts the results from ONETEP and is correctly rectified in the extended catalogue.

Of course, restricting the catalogues to pairs is arbitrary: like in diagrammatic theories, we could include triplet, quadruplets, *et cetera*. To gauge how our approximation is performing, we have simulated a system with 4096 atoms and 29 S impurities, both with our effective model and ONETEP. As shown in Fig. 3.4, we obtain a spectrum of impurity states that extends roughly over the same range of energies, especially towards the valence band. The number of states in the spectrum also matches between the effective model and the ONETEP simulation, i.e. we obtain the correct number of bonding and anti-bonding states, as discussed in the previous paragraph. We note, however, that the effective model slightly underestimates the energy in the upper portion of the spectrum, resulting in a larger band gap than computed by the DFT, where the impurity and conduction band seem to have already joined.

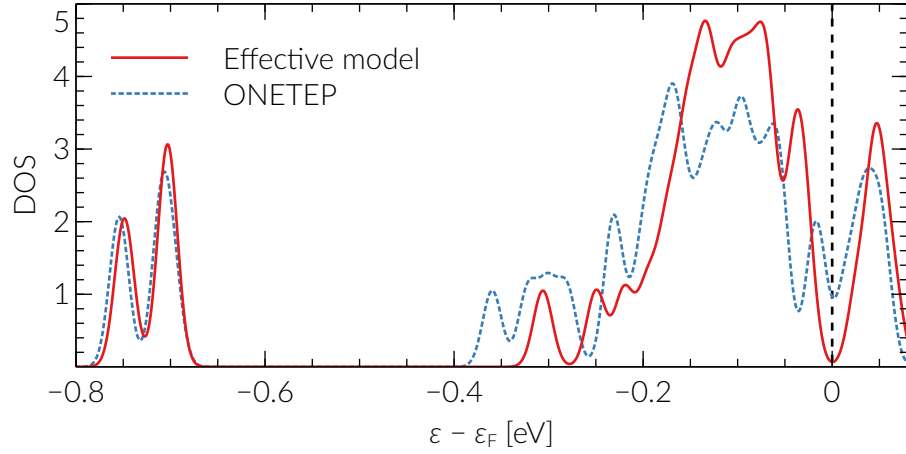


Figure 3.4: Density of states of an exemplary system of 4067 Si atoms and 29 S impurities, obtained by summing Gaussian distributions with $\sigma = 10$ meV centred over the eigenvalues. We show the spectrum obtained from the ONETEP simulation with a blue dashed line, while that from its effective model is in red continuous. The spectra are aligned at the Fermi energy ε_F (dashed line), defined as the midpoint between the energies of the highest occupied and lowest unoccupied states. Reproduced from the Supplemental Materials to Carnio et al. (2017).

3.3 Technical construction

We now delve into the technical aspects of the construction of the effective models. The starting point is the list of atom positions, with the indication of their species. The format we use is simply that of a ONETEP input file, Fig. 2.1, except that coordinates are expressed in units of lattice parameter a . For the specific case of Si, we can actually re-express all coordinates in units of $a/4$, in order to work with integer coordinates. One might avoid this “post-processing” and use integer coordinates consistently from the very beginning. In this case we do not need to take into account the coordinates of the *relaxed* lattice: the displacement of the lattice sites (Fig. 2.7) occurs on a much smaller length scale than the interatomic separation, hence we can always uniquely map the sites of the original and the deformed lattices.

3.3.1 Description in map

The list of coordinates is subsequently converted to a *map* of the system, exemplified in Fig. 3.5. For each atom we report, in order:

1	-1	1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
2	1	2	0.75	0.25	0.25	-0.75	-0.25	1.25	0.00	0.00	0.00
3	2	2	0.25	0.75	0.25	0.00	-1.50	-0.50	0.25	-1.75	-0.25
						.	.	.			
12	3	2	0.75	0.25	1.25	-0.75	0.75	-0.75	-0.75	1.25	-1.25
						.	.	.			
123	0	2	2.75	2.25	2.25	0.00	0.00	0.00	0.00	0.00	0.00
						.	.	.			

Figure 3.5: Example of a map of the doped Si system.

index The index i in the list.

catalogue The catalogue c , namely the source of the elements in the Hamiltonian used to describe this site. It is an integer that can have values -1 (the site is a Si atom more than 10 shells away from any impurity), 0 (the site is occupied by an impurity), 1 (a Si atom within 10 shells from an isolated impurity), and 2 to 5 (a Si within 10 shells from a pair of defects).

sublattice An index s indicating whether the site belongs to sublattice ‘1’ or ‘2’ as defined in (2.12).

position The position \mathbf{r}_i of site i .

distances to closest defects The displacement vector(s) $\delta_{i,1}$ ($\delta_{i,2}$) to the closest (second closest) defect, or 0 if $c < 1$ ($c < 2$).

This description contains all the information needed to build our effective systems. It is definitely not the most efficient: the catalogue c , for instance, can be deduced from $\|\delta_{i,1} - \delta_{i,2}\|$. Nine Cartesian coordinates are also redundant, when we could use polar angles and triangulation. Notice, however, that such description might involve real numbers with a finite-precision representation. This might create problems when reading from and comparing to strings of numbers from the catalogues. A redundant description with integer coordinates, as described above, makes this process easier and more error-proof.

3.3.2 Creation of the catalogues

Now that we have introduced the map, we can discuss how we build the catalogues from the ONETEP-simulated prototypes, starting from those with dopants.

From the map of a prototype we know where the impurities ($c = 0$) and the surrounding shells ($c \geq 1$) are. We shall remind the reader here that our matrices

are divided in 9×9 blocks, hence, for example, the impurity will be described by a “row of blocks” i_0 (we will call it “block-row” from now on). From this block-row we are interested in the block of self-energies (on the block-diagonal) and the hopping terms to the neighbours. If one such block of hopping terms is j , in the catalogue we will store

$$0, s, \mathbf{r}_j - \mathbf{r}_{i_0}, (0, 0, 0), (0, 0, 0), \quad (3.7)$$

together with the 81 real numbers forming the block. By changing j , the displacement $\mathbf{r}_j - \mathbf{r}_{i_0}$ will point to all neighbours of site i_0 and, as stated before, we only store those up to the 10th shell.

This example shows the general procedure to build the catalogue: we start from the map, single out all elements with $c \geq 0$ and, for each of them, store self-energies and hopping blocks to their neighbours. To keep things in order, we create a different file for each value of c and for Hamiltonian or overlap matrix elements.

The gathering of the catalogues and its subsequent reuse in building the effective models need to be done *consistently*, i.e. there are a number of details that need to be taken care of, or things *will* go wrong. For what concerns the catalogues:

- we need to symmetrize the source matrix, which might be not symmetric due to numerical fluctuations introduced by the DFT;
- if for a site i we have $\|\delta_{i,1}\| = \|\delta_{i,2}\|$, so the two closest impurities are at the same distance, we store catalogue entries both in the order $\delta_{i,1}, \delta_{i,2}$ and $\delta_{i,2}, \delta_{i,1}$.

We conclude with some words on the construction of $c = -1$, the matrix elements to describe the host material. In this case, from a prototype of 4096 Si atoms, we have 2048 sites for each sublattice s . For each of these sites, the self-energy and hopping terms should be equal, up to numerical fluctuations of around 0.05% introduced by the DFT simulation. We counteract this effect by averaging over all blocks with the same s and displacement vector.

3.3.3 Construction of the effective matrices

To build the effective models we first load in the memory the map of the system and the catalogues. Most of the work, then, is done by a double loop that computes

the catalogue entry connecting two sites and looks for the corresponding matrix elements in the catalogues.

More specifically, we only build the upper triangles of the matrices, since they are symmetric. We therefore scan through the 9×9 blocks using indices $i = 1, \dots, N$ (for the block-rows) and $j = i, \dots, N$ (block-columns), where N is again the number of sites. If we want to fill in block (i, j) , we retrieve the catalogues c_i and c_j that describe block-rows i and j . Notice that we are still talking about rows, here, since the catalogue is stored in rows.

Let us take the case where $c_i = -1$ and $c_j = 1$: when we scan through the catalogue for i we will be looking for hopping terms to a site that is influenced by an isolated defect. But such terms do not exist in the $c = -1$ catalogue! For this reason, we have to introduce the check $c_i < c_j$, which, if true, triggers the swapping of i with j . When we do so, however, we will retrieve in the catalogue the hopping terms $j \rightarrow i$, while we want those $i \rightarrow j$, which means that the matrix elements we find will need to be transposed with respect to the main diagonal of the matrix. Without the implementation of this check, the matrices we build will be incorrect. Unfortunately, because we feed the upper-triangular matrices to the diagonalization routine, we would still obtain real (but wrong!) eigenvalues.

Another detail that needs attention is the implementation of *periodic boundary conditions*, both in the construction of the catalogues and of the effective models. The shells of neighbours for a site might extend into the periodic copies of the system and these points must be mapped into the original system. In fact, we identify such neighbours by looping over all points in the system and checking if their equivalent site in a periodic copy is within the shells cut off. This operation usually requires modulo operations to map a distance $\delta = x_i - x_j$ in the interval $[-L_x/2, L_x/2]$. If the box size in the x direction is L_x , in Mathematica (Wolfram Research Inc., 2015) we will write `Mod[xi-xj, Lx, Lx/2]`, whereas in Fortran we will compute

$$\tilde{\delta} = \begin{cases} (\delta + L_x/2 \bmod L_x) - L_x/2 & \delta \geq 0 \\ (\delta - L_x/2 \bmod -L_x) + L_x/2 & \delta < 0 \end{cases}. \quad (3.8)$$

We have provided these examples because the definition of the modulo operation, especially when the dividend is negative, can change between programming languages.

Finally, depending on the diagonalization routine, the output will have to be

cast in a specific format. We save space by storing the matrices in the Compressed Sparse Row format (Saad, 2011), i.e. we ignore the zero values. Moreover, when it comes to the blocks on the diagonal, we shall remember to discard the values in the lower triangle. The sparsity of the Hamiltonian (overlap) matrix is $662/(9N)$ ($554/(9N)$), where the specific numbers depend on the number of neighbours included in the catalogue, i.e. on the “near-sightedness” of our DFT simulations (see Sec. 3.2).

3.4 Diagonalisation

Like for the Kohn-Sham Hamiltonian (2.7), once we have the Hamiltonian H and overlap S matrices we diagonalize them to find eigenvalues ε_j and eigenvectors ψ_j :

$$H |\psi_j\rangle = \varepsilon_j S |\psi_j\rangle \quad j = 1, \dots, 9N^3. \quad (3.9)$$

The presence of S , also known as *mass matrix*, makes Eq. (3.9) a *generalised eigenvalue problem*. Because both matrices are symmetric and S is positive-definite by Eq. (2.9), the eigenvalues are real and the corresponding eigenvectors are mutually orthogonal with respect to the scalar product defined by S : $\langle \psi_i | S \psi_j \rangle = \delta_{i,j}$ (Saad, 2011). If this does not happen, something has gone wrong in the implementation discussed in the previous Section.

The eigenvectors that solve (3.9) are written in the basis of support functions discussed in Sec. 2.1.5, i.e. in the basis with 9 orbitals per site: $|\psi_j\rangle = \sum_{\alpha} M_{\alpha}^j |\phi_{\alpha}\rangle$. For the analysis that follows, we write the eigenvectors in a “site” basis by summing over the nine orbital coefficients of each site k :

$$|\Psi_k^j|^2 = \sum_{\alpha \in k} \sum_{\beta} M_{\alpha}^j S_{\alpha\beta} M_{\beta}^j. \quad (3.10)$$

This transformation also deals with the unavoidable consequences of working with a non-orthogonal basis: while $\sum_k |\Psi_k^j|^2 = \langle \psi_j | S \psi_j \rangle = 1$, we might have $M_{\alpha}^j \leq 0$ for some orbital α . If after the summation we have $|\Psi_k^j|^2 \geq 0, \forall k$, we interpret (3.10) as the probability of finding an electron in site k .

3.5 Numerical challenges and the choice of sulphur

We tackle Eq. (3.9) numerically using JADAMILU (Bollhöfer and Notay, 2007), a code to obtain few eigenvalues and eigenvectors of large sparse matrices. This library has already been employed to diagonalise the 3D Anderson model (AM, Schenk et al. (2006)) in different numerical studies (Rodriguez et al., 2009, 2011; Thiem and Schreiber, 2013; Harashima and Slevin, 2014; Puschmann et al., 2015; Moore et al., 2017).

Let us imagine an effective system with N atoms in total, N_X of which are impurities. The corresponding concentration n is given by

$$n = \frac{N_X}{\text{Volume}} = \frac{N_X}{N} \frac{8}{a^3}, \quad (3.11)$$

where a is the lattice constant of Si. If we want to study the impurity states induced in the band gap we need to ask JADAMILU for at least N_X eigenvectors, around a certain target ε_t . From Eq. (3.11), however, we can see that N_X , at a given concentration, increases with N . In comparison, the spectrum of the AM is symmetric and the delocalisation-localisation transition is studied only for the state with energy closest to 0.

The features that make our model realistic, compared to the AM, are also those that make it more complex and difficult to treat numerically. First of all, since our matrices describe the 9 orbitals as well, the size of our matrices is 9 times larger. Moreover, we allow hopping up to the 10th shell in the diamond cubic lattice, for a total of 146 neighbours; in comparison, the hopping in the AM is only to the nearest neighbours, 6 for the simple cubic lattice. As mentioned in the previous section, the larger number of neighbours in our model translates in Hamiltonian (overlap) matrices that have 662 (554) non-zero elements per row, two orders of magnitude more than the 7 in the AM.

Nevertheless, the construction and diagonalization of a 4096-atom system with a concentration of S dopants of 1% takes around half an hour on a recent Intel processor, which is considerably less than the 12 h on 1152 cores needed for the DFT simulations of the prototypes. It is also considerably more than what is required to simulate an AM with the same number of sites.

The additional complexity in our model also prevents us from diagonalizing systems as large as those in the recent studies of the 3D AM (Rodriguez et al., 2011; Lindinger and Rodríguez, 2017). For this reason, the results presented in Ch. 4, 5




Figure 3.6: Dependence of the concentration on the number of impurities N_x and of total atoms N , from Eq. (3.11). As an example we show the smallest and largest system sizes we consider in this work, namely $N = 16^3$ (red) and $N = 22^3$ (blue), as well as $N = 30^3$ (green) for comparison. Horizontal lines indicate the critical concentrations for Si:P (violet) at approximately $3.5 \times 10^{18} \text{ cm}^{-3}$, and for Si:S (orange) at $2 \times 10^{20} \text{ cm}^{-3}$. The latter value was arbitrarily chosen knowing that it falls between 1.8 and $4.3 \times 10^{20} \text{ cm}^{-3}$ (Winkler et al., 2011).

and 6 correspond to Si:S, only. As mentioned in Sec. 2.4, the critical concentration for Si:P ($3.52 \times 10^{18} \text{ cm}^{-3}$) is two orders of magnitude smaller than that of Si:S ($\sim 10^{20} \text{ cm}^{-3}$). This implies that simulating a system with N_x P impurities requires two orders of magnitude more host atoms than in the case of S dopants. In fact, if we consider a system with 10 648 atoms, which is the largest we can diagonalize, a single P impurity corresponds to a concentration $n = 4.7 \times 10^{18} \text{ cm}^{-3}$, greater than the critical value and hence already in the metallic regime. As shown in Fig. 3.6, we need much larger systems, for instance $N = 30^3 = 27\,000$, before we have just a couple of states in the localised regime.

Chapter 4

Multifractality of the wave functions

After diagonalising the effective models and obtaining the impurity band states, we are going to analyse their statistical properties. From the theory of critical phenomena, briefly reviewed following Cardy (1996) in Sec. 4.1, we retrace how the critical wave function at the Anderson transition was understood to have a multifractal nature.

With this motivation, in Sec. 4.2 we follow Tél (1988) and Janssen (1994) in reviewing the basic concepts and definitions behind multifractals. These concepts are then translated in the context of disordered physics and used to analyse the statistical properties of the wave functions from our simulations.

The reader should be aware that it is difficult to disentangle the results from this chapter to those that follow. In particular, the multifractal analysis relies, at least for what this chapter concerns, on the knowledge of the critical point of the Anderson transition, which is estimated in Ch. 5.

4.1 Continuous phase transitions

A *phase transition* is a sudden change in the properties of a system (a *phase*) caused by the variation in external conditions. These changes are reflected in the non-analytical behaviour of thermodynamic quantities at specific points in parameter space, the *critical points*. In the case of the Anderson transition, an increase in the degree of disorder drives the system from a delocalised (or extended) phase to a localised phase (see Sec. 3.1).

Phase transitions are called *discontinuous*, or *first-order*, if the relevant thermodynamic quantities show discontinuities when passing through the critical point. In a *continuous* or *second-order* phase transition, instead, the divergence is seen in the correlation length ξ , i.e. the distance over which the fluctuations in the microscopic degrees of freedom of the material are correlated. At criticality, then, we observe the *absence of length scales*: the fluctuations in the system are correlated at all length scales, at least down to microscopic distances like the lattice constant a . The system must be in a unique (critical) phase that is continuously connected to the phases on each side of the transition, implying that the thermodynamic quantities change smoothly across the transition.

For the Anderson transition, the absence of length scales means that the wave function at the critical disorder is *self-similar* (Aoki, 1982). It also needs to have a “filamentary” structure (Aoki, 1983), i.e. it needs to be extended throughout the volume, but also occupy an infinitesimal fraction of it, a property of the localised phase. This structure, then, allows the critical phase to be continuously connected to both the extended and localised phases. In conjunction with its self-similar property, the critical wave function qualifies as a *fractal*, at least if we disregard the lower limit imposed by a . A few years later, Castellani and Peliti (1986) realised, based on the earlier work of Wegner (1980), that the critical wave function is not a simple fractal, but rather an “interwoven family” of fractals, each with its own dimension and distribution. Such an object is a *multifractal* (Mandelbrot, 1972, 1974).

4.2 Multifractals

To better understand the link between self-similarity and fractality, let us consider a system occupying a finite region of space $\Sigma \subset \mathbb{R}^D$ with a local density $\rho(\mathbf{r})$. Following Pietronero (1990), we define the pair correlation function,

$$g(\mathbf{r}) = \langle \rho(\mathbf{r} + \mathbf{r}') \rho(\mathbf{r}') \rangle_{\mathbf{r}'}, \quad (4.1)$$

which gives the probability that two points separated by \mathbf{r} both belong to the region Σ . For simplicity, we now assume that the correlation function is isotropic, $g(\mathbf{r}) = g(r)$. In the absence of length scales, g obeys *homogeneity laws* (or *scale-invariance*) with respect to a *resolution* or *coarse-graining* λ . More specifically, if we

rescale lengths as $r \rightarrow r' = \lambda r$ we have that

$$g(r') = \lambda^\kappa g(r), \quad (4.2)$$

where κ is a *homogeneity exponent*. The solution to this equation is given by a *power-law behaviour*, $g(r) \propto r^\kappa$. If we then fix r as the reference length scale, and $g(r) = 1$,

$$g(\lambda) = \lambda^\kappa \quad (4.3)$$

translates self-similarity into a mathematical relation.

We can finally define fractal objects as self-similar structures whose observed spatial extent (e.g. volume) depends, with a power-law behaviour, on the resolution at which we look at it. For fractals originating from a mathematical relation, the dependence on the resolution can extend over an infinite range. For fractals appearing in physical systems, instead, the range of λ is usually limited by macro- and/or microscopic scales. A very comprehensive list of examples can be found in Malcai et al. (1997).

4.2.1 Measures, fractals and multifractals

Let $\psi(\mathbf{r})$ be the wave function of an electron in a $L \times L \times L$ volume. The modulus square $|\psi(\mathbf{r})|^2$ defines a normalised measure on this volume, which we partition in boxes of linear size $l = \lambda L$. The number of boxes will then be λ^{-d} , where $d = 3$ is the Euclidean dimension of the support of the system. The probability of finding the electron in box i , then, is the *box-probability*

$$\mu_i = \int_{\text{box } i} |\psi(\mathbf{r})|^2 d\mathbf{r}. \quad (4.4)$$

We can then define the *fractal dimension* D of the system by counting the number of boxes where the box-probability does not vanish:¹ $N(\lambda) \sim \lambda^{-D}$. Because the electron can access any portion of the volume, i.e. there is no region of space with vanishing probability, we conclude that the fractal dimension $D = d = 3$, which is not very interesting.

¹It is customary to use \sim to indicate that the proportionality constant is independent of the resolution and can thus be ignored. This constant might appear, for instance, when the boxes, whichever their shape, do not perfectly cover the system. Since we are covering boxes with boxes, most of the relations in this section are actually equalities.

Compared to the fractal dimension, more insightful is actually the study of the powers of the box-probability μ_i^q , which is the idea behind *multifractal* analysis. If the wave function is a multifractal, we expect to see the power-law behaviour of (4.3):

$$\langle \mu^q \rangle_L \sim \lambda^{D+\tau_q}, \quad (4.5)$$

where $\langle \dots \rangle_L$ denotes the average over all boxes in the volume. Equivalently, we can introduce the *partition sum* $R_q(\lambda)$ (also the *generalised inverse participation ratio*) as $\langle \mu^q \rangle_L = \lambda^D R_q(\lambda)$ and write

$$R_q(\lambda) = \sum_i \mu_i^q \sim \lambda^{\tau_q}. \quad (4.6)$$

The *mass exponents* τ_q describe the scaling behaviour of the moments and do not depend on λ .

Let us stress again that multifractality holds if, in the power-law relation of Eq. (4.5), $\tau_q \neq 0$ for a finite range of λ : the box size l should be smaller than the system size, but also larger than the microscopic scale a . At the same time, for critical states at the Anderson transition, the system size is much smaller than the correlation length ξ , such that

$$a \ll l < L \ll \xi \quad (4.7)$$

Additionally, the wave function is a truly critical (and hence multifractal) only in the thermodynamic limit, where both L and ξ diverge, hence τ_q is uniquely defined in the limit $\lambda \rightarrow 0$. For finite systems, instead, we choose states and coarse-grainings that satisfy (4.7). In this case, we can *estimate* τ_q by fitting the slope of $\log R_q(\lambda)$ versus $\log \lambda$. We are assuming here that multifractality survives in finite systems (Cuevas and Kravtsov, 2007), and postpone the discussion of this non-trivial assumption to Ch. 5.

From (4.5) and the normalisation of the wave function, it is possible to show that $\tau_0 = -D$ and $\tau_1 = 0$. This implies that we can generalise the definition of the fractal dimension to a function D_q such that $D_0 = D$ and $\tau_q = D_q(q-1)$. In the case of a simple fractal $D_q \equiv D$, while for a multifractal D_q has a non-trivial dependence on q . The deviation from the simple-fractal case is captured by the *anomalous scaling exponent* $\Delta_q = (D_q - d)(q-1) = \tau_q - d(q-1)$.

4.2.2 The multifractal spectrum

The scaling of the moments R_q , yielding τ_q , is enough to fully characterise the multifractal nature of the wave function. Now we present an equivalent description of the multifractal that will be useful, in the sections that follow, to validate our results and compare them to the 3D Anderson model. This description is founded on a *multifractal measure* (Frisch and Parisi, 1985), a distribution such that, around each box i , $\mu_i = \lambda^{\alpha_i}$. The set of boxes with $\alpha_i \in [\alpha, \alpha + d\alpha]$, then, constitutes a simple fractal with dimension $f(\alpha)$, such that the number of said boxes is

$$N_\lambda(\alpha) \sim \lambda^{-f(\alpha)} \quad \text{and} \quad \alpha_i = \frac{\log \mu_i}{\log \lambda}. \quad (4.8)$$

This is the formalisation of the idea of Castellani and Peliti (1986), where the multifractal is composed of different simple fractals.

We re-express the partition sum of Eq. (4.6) as

$$R_q(\lambda) = \sum_i \mu_i^q = \sum_i \lambda^{q\alpha_i} = \int N(\alpha) \lambda^{q\alpha} d\alpha \sim \int \lambda^{q\alpha - f(\alpha)} d\alpha. \quad (4.9)$$

For small λ , we can use the saddle point approximation and find that the biggest contribution in the integral (4.9) comes from the value of α that maximises (since $\lambda < 1$) the argument of the exponential, i.e. the α_q such that $f'(\alpha_q) = q$. We can then write, from (4.6), $\tau_q = q\alpha_q - f(\alpha_q)$. If we identify $f_q = f(\alpha_q)$ we can see that (q, τ_q) and (α_q, f_q) are related by a Legendre transformation:

$$f_q = q\alpha_q - \tau_q \quad \text{and} \quad \alpha_q = \frac{d\tau_q}{dq}. \quad (4.10)$$

It can be proven, e.g. in Janssen (1994), that τ_q is a monotonically increasing function in q , which implies that $\alpha_q > 0, \forall q$.

We can combine *singularity strengths* α_q and the *singularity spectrum* f_q to obtain the *multifractal spectrum* $f(\alpha)$. This function is equivalent to the generalised dimensions D_q in characterising the multifractal, and in the case of a simple fractal analogously reduces to the point (D, D) in a $(\alpha, f(\alpha))$ plot. As shown in the example of Fig. 4.1, $f(\alpha)$ is a convex function reaching its maximum at α_0 with a value $f_0 = \tau_0 = D$. From (4.10) we further notice that $f_1 = \alpha_1$, since $\tau_1 = 0$. The spectrum is therefore tangential to the functions $f_0(\alpha) \equiv D$ and $f_1(\alpha) = \alpha$.




Figure 4.1: Multifractal spectrum $f(\alpha)$ for the ONETEP prototype described in Ch. 2, computed for q from -2 to 5 in steps of 0.1 (increasing from right to left). Dashed lines indicate the functions $f_0 \equiv D$ and $f_1(\alpha) = \alpha$.

4.2.3 Symmetry of the multifractal spectrum

Using the nonlinear σ model, Mirlin et al. (2006) have analytically proven that at criticality the multifractal exponents (4.10) satisfy the exact symmetry relation

$$\alpha_q + \alpha_{1-q} = 2d \quad f_{1-q} = f_q + d - \alpha_q . \quad (4.11)$$

Assuming the universality of the critical properties at the Anderson transition, this result is expected to generally hold for the Wigner-Dyson symmetry classes (see Ch. 1). Indeed, this result was confirmed numerically for different systems, including the 3D Anderson model (Rodriguez et al., 2008; Vasquez et al., 2008) and experiments (Faez et al., 2009)

4.3 Multifractal analysis of the wave function

In the rest of the chapter we are interested in the singularity strengths α_q , which, together with the mass strengths τ_q and the anomalous dimensions Δ_q are called multifractal exponents (MFE). In this section we recast the exponents derived in Sec. 4.2 in a form that is more convenient for numerical calculations, mostly by

reducing the loss of precision. We then extend our definitions to include a disorder average.

4.3.1 Numerical calculation

Following Chhabra et al. (1989), it is convenient to define, from (4.6) and (4.10), the auxiliary quantity

$$S_q(\lambda) = \frac{d R_q(\lambda)}{d q} = \sum_i \mu_i^q \log \mu_i . \quad (4.12)$$

This ratio can be interpreted as an average with respect to the measure defined by μ^q . The latter is also called q -microscope, because it increases the large (small) box-probabilities for $q > 0$ ($q < 0$). A computationally-friendly formulation of the MFE reads

$$\tau_q = \lim_{\lambda \rightarrow 0} \frac{\log R_q(\lambda)}{\log \lambda} \quad \Delta_q = \tau_q - q(d-1) \quad \alpha_q = \lim_{\lambda \rightarrow 0} \frac{S_q(\lambda)}{R_q(\lambda) \log \lambda} . \quad (4.13)$$

To comply with (4.7), we choose $\lambda \leq 1/2$, namely we consider boxes of linear size up to $l \leq L/2$. We coarse-grain the wave function using the partitioning scheme proposed by Thiem and Schreiber (2013). Here, the box size l can take any integer value (up to $L/2$), so that $\lambda^{-1} = L/l$ can take non-integer values. This is achieved by first periodically replicating the original system, such that it can be exactly covered by an integer number of boxes, and then by averaging over the possible equivalent box origins. The increased number of available box sizes translates, in the linear fits, in reduced uncertainties in the estimated slopes.

4.3.2 Ensemble averaging

So far we have computed the multifractal properties of a single wave function. The multifractal analysis of the Anderson transition is usually performed by taking an average over the disorder realisations. The definitions of the MFE can be extended by defining the *ensemble average* of the partition sum as $\langle R_q(\lambda) \rangle \sim \lambda^{\tau_q^{\text{ens}}}$, such that

$$\tau_q^{\text{ens}} = \lim_{\lambda \rightarrow 0} \frac{\log \langle R_q(\lambda) \rangle}{\log \lambda} . \quad (4.14)$$

We then proceed to take the Legendre transform and define

$$f_q^{\text{ens}} = q\alpha_q^{\text{ens}} - \tau_q^{\text{ens}} \quad \text{and} \quad \alpha_q^{\text{ens}} = \frac{d\tau_q^{\text{ens}}}{dq} = \lim_{\lambda \rightarrow 0} \frac{\langle S_q(\lambda) \rangle}{\langle R_q(\lambda) \rangle \log \lambda}. \quad (4.15)$$

Notice that, in the ensemble average of α_q , the q -microscope μ_i^q is normalised by $\langle R_q \rangle$, namely the averaged partition sum of the wave function. If we normalised the μ_i^q terms for every wave function we would obtain the *typical* average $\alpha_q^{\text{typ}} = \langle S_q/R_q \rangle / \log \lambda$ ($\lambda \rightarrow 0$). While in the ensemble average all wave functions, including rare events, are equally weighted, the typical average is dominated by the behaviour of “typical” wave functions.² The presence of rare events translates in the appearance of negative fractal dimensions (see Sec. 4.5.1), a feature of the $f(\alpha)$ that is best captured by ensemble averaging. Please consult Evers and Mirlin (2008), and references therein, for more details.

The standard deviation σ_{α_q} associated to α_q^{ens} (at a fixed λ) is related to the standard deviations σ_{S_q} and σ_{R_q} and the covariance $\text{cov}(S_q, R_q)$ via propagation of the variance (Rodriguez et al., 2011):

$$\sigma_{\alpha_q} = \alpha_q^{\text{ens}} \sqrt{\frac{\sigma_{S_q}^2}{\langle S_q \rangle^2} + \frac{\sigma_{R_q}^2}{\langle R_q \rangle^2} - 2 \frac{\text{cov}(S_q, R_q)}{\langle S_q \rangle \langle R_q \rangle}}, \quad (4.16)$$

and, analogously,

$$\sigma_{\tau_q} = \frac{\sigma_{R_q}}{\langle R_q \rangle \log \lambda}, \quad (4.17)$$

so that, finally, $\sigma_{f_q}^2 = \sigma_{\alpha_q}^2 + \sigma_{\tau_q}^2$. The standard error of the mean is obtained by dividing every standard deviation by $\sqrt{\mathcal{N}}$, where \mathcal{N} is the number of available realisations.

For finite systems, the multifractal exponents are computed, as explained in Sec. 4.2.1, by estimating the slope of a $\log \langle R_q(\lambda) \rangle$ vs. $\log \lambda$ plot, in the case of τ_q^{ens} . Accordingly, the uncertainties (4.16) and (4.17) on the data points at fixed λ have to be multiplied by a factor $\log \lambda$.

4.4 Statistics and disorder average

Before applying the multifractal analysis to the wave functions calculated in the effective model, let us now resume the discussion of chapters 2 and 3. For each

²This is the same difference between arithmetic and geometric mean.

system size $N = L^3$ we diagonalize realizations with N_S randomly-distributed impurities, with values chosen to regularly sample approximately the same range of concentrations n , given by Eq. (3.11). In Table 4.1 we report the number of realizations \mathcal{N} computed for each system size, and the total number of corresponding wave functions.

For each realisation we compute $N_S + 5$ states, to ascertain that we obtain a complete impurity band (IB). The additional states usually come from the conduction band (CB), but, especially with very few impurities, we have also found states from the valence band (VB). This requires criteria to distinguish such bands. The VB states, assuming there are N_{VB} of them, are filtered out by checking that (i) their energy is less than 6.28 eV (the top of the VB in the empty Si system), and (ii) their participation ratio³ is greater than the convenient value 0.3. This value is a safe lower bound that we have chosen after studying a handful of realisations which included VB states. From the $N_S + 5 - N_{VB}$ states left, the lower $N_{IB} = N_S - N_B$ are impurity band states, where N_B is the number of bonding states that have disappeared into the VB (cf. Sec. 3.2). The remaining states are then CB states.

After categorising the states, we can define the Fermi energy ε_F as the midpoint between the energy of the highest IB state, which is occupied when we build the many-body wave function of all electrons, and the energy of the lowest (unoccupied) CB state.

We then redefine $\varepsilon_F = 0$ for each realization and partition the negative energies in several bins centred on $\{E_i\}_i$. For each bin we consider the eigenstate Ψ^j (in the site basis, see 3.4) whose energy ε_j is closest to E_i , and calculate the auxiliary quantities (4.6) and (4.12), which we include in the ensemble average of the multifractal exponents (4.15).

4.5 Results from the effective models

In Fig. 4.2 we show the average singularity spectrum for the ensemble of 10 648 atoms with 140 impurities, a system that is critical at energy close to -0.249 eV (estimated in Ch. 5). The increase in the ensemble size does not change the shape of the spectrum significantly, but has the effect of reducing the error bars on the data points. In particular, it is the calculation of the average $\langle R_q \rangle$ and $\langle S_q \rangle$ in Eq.

³From Eq. (4.6) calculated as $[NR_2(1/L)]^{-1}$.

N	N_S	n [10^{20} cm^{-3}]	$\langle \mathcal{N} \rangle$	$(\mathcal{N}_{\min}, \mathcal{N}_{\max})$	$\sum \mathcal{N}$	Eigenstates
16^3	4–200	0.49–24	802	(200,1000)	68 153	2 943 811
18^3	5–322	0.43–28	758	(106,1000)	64 430	3 951 351
20^3	5–365	0.31–23	732	(162,1000)	71 051	5 640 229
22^3	10–410	0.47–19	541	(293,733)	34 067	5 521 425
Total no. of realizations and wave functions:					237 311	18 056 816

Table 4.1: Summary of the range of impurities N_S , the concentration n , the average, minimum, maximum and total number of disorder realizations for each L , indicated by $\langle \mathcal{N} \rangle$, \mathcal{N}_{\min} , \mathcal{N}_{\max} and $\sum \mathcal{N}$, respectively. The final column indicates the total number of eigenstates calculated per system size, and the last row the total for all $N = L^3$ and n . Reproduced from Carnio et al. (2017).

(4.15) that benefits from larger ensembles, since smaller error bars in the data used for the linear fits implies smaller uncertainties on the fit parameters, see Fig. 4.3.

We quantitatively report the quality of said fits in the lower panel of Fig. 4.2, where we show the linear correlation coefficient r^2 and the p value. As noted in Rodriguez et al. (2008), while $r^2 \approx 1$ indicates a good linear behaviour, small p values suggest that the uncertainties on the data point are too small to support the deviation from the linear behaviour we are fitting. This is likely due to the limited number of realisations available for the ensemble averaging. For comparison, at the end of the two branches, i.e. for large $|q|$ values, error bars are larger and hence the quality-of-fit increases again.

4.5.1 Negative fractal dimensions

Error bars increase on the two ends of the spectrum, for negative q (right) and positive (left). For $q < 0$, the q -microscope increases the weight of small values of the wave function, which are more sensitive to numerical fluctuations from the diagonalization. The other end of the spectrum ($q > 0$, left) describes instead the presence of rare critical functions with small values of α and hence large $|\psi|^2 \sim L^{-\alpha}$. The set of these values scales with a negative fractal dimension $f(\alpha)$, which means that their occurrence frequency vanishes in the $L \rightarrow \infty$ limit. This is a known effect arising from ensemble averaging (Chhabra and Sreenivasan, 1991) and known since the pioneering work of Mandelbrot (Mandelbrot, 1984). This finite-size effect is further observed and commented in Rodriguez et al. (2008), where larger systems are accessible and studied. In comparison, the single state used to produce Fig. 4.1 does not show said rare boxes with large probability

amplitudes, as indicated by $f_q > 0$.

4.5.2 Width of the multifractal spectrum

Finally, let us comment on the width of the distribution in Fig. 4.2, as compared to the Anderson model studied in Rodriguez et al. (2008). A narrow $f(\alpha)$ spectrum implies that extreme values (either large or small) occur less frequently. This means that, in our case, the average state near criticality in our model looks more homogeneous or extended than in the Anderson model. In Fig. 4.4, instead, we show the singularity spectrum for a system of 10 648 atoms, which, for 230 impurities, is close to criticality at the Fermi energy $\varepsilon_F = 0$ and deeper in the impurity band at -0.320 eV (estimated in Ch. 5). While both spectra are narrower than the Anderson model, the critical wave function at ε_F appears on average more extended than deeper in the impurity band. This observation will resurface again in Chapters 5 and 6.

Our results are reminiscent of those found by Mirlin and Evers (2000) for the power-law random banded matrix (PRBM) model, which describes a 1D chain with random long-range hopping decaying as $r^{-\alpha}$ over distances larger than a band width b . For the critical value $\alpha = 1$, the model undergoes an Anderson transition for any value of b , which parametrises a family of critical models that can be studied from the weak- ($b \gg 1$) to the strong-coupling ($b \ll 1$) regime. For $b \gg 1$ the model shows a “quasi-metallic” behaviour, where the critical wave functions show statistical properties similar to the delocalised phase. The singularity spectrum becomes correspondingly narrower with a parabolic shape, a regime called *weak multifractality*. In this case the multifractal spectrum follows the parabolic *approximation* (Janssen, 1994):

$$f(\alpha) \simeq d - \frac{(\alpha - \alpha_0)^2}{4(\alpha_0 - d)} \quad \text{and} \quad \alpha_0 = d + \gamma. \quad (4.18)$$

In Fig. 4.2 we fit our full-ensemble data to (4.18) to find the estimate value $\alpha_0 \approx 3.55$. We only report an approximate value without uncertainty because the parabolic behaviour is an approximation and does not necessarily hold for the whole spectrum, since $f(\alpha)$ is defined only for positive α . In fact, at $q_c = (d + \gamma)/2\gamma$ (Evers and Mirlin, 2008) we have $\alpha_{q_c} = 0$ and $f(0)$ is finite (a *termination point*), whereas our results in Fig. 4.2 and those in Rodriguez et al. (2008) seem to suggest that $f(\alpha) \rightarrow -\infty$ in the limit $\alpha \rightarrow 0$ (hence no termination point).

4.5.3 Symmetry of the multifractal spectrum

In Fig. 4.2 we also show the symmetrised spectrum obtained by computing and plotting $\{\alpha_{1-q}, f_{1-q}\}$ from Eq. (4.11). As expected from the previous paragraphs, the uncertainty on the data points increase at the extremities. Within these error bars, the spectra are in good agreement with each other. We verify the same symmetry relation also for the spectra in Fig. 4.4. While at -0.320 eV there is excellent agreement between the spectra, at the Fermi energy there is a slightly higher discrepancy, especially at extreme values of q . Since this discrepancy would be resolved by taking two standard deviations as confidence intervals, instead of one, we cannot attribute this discrepancy to any specific underlying physical factor or systematic error.

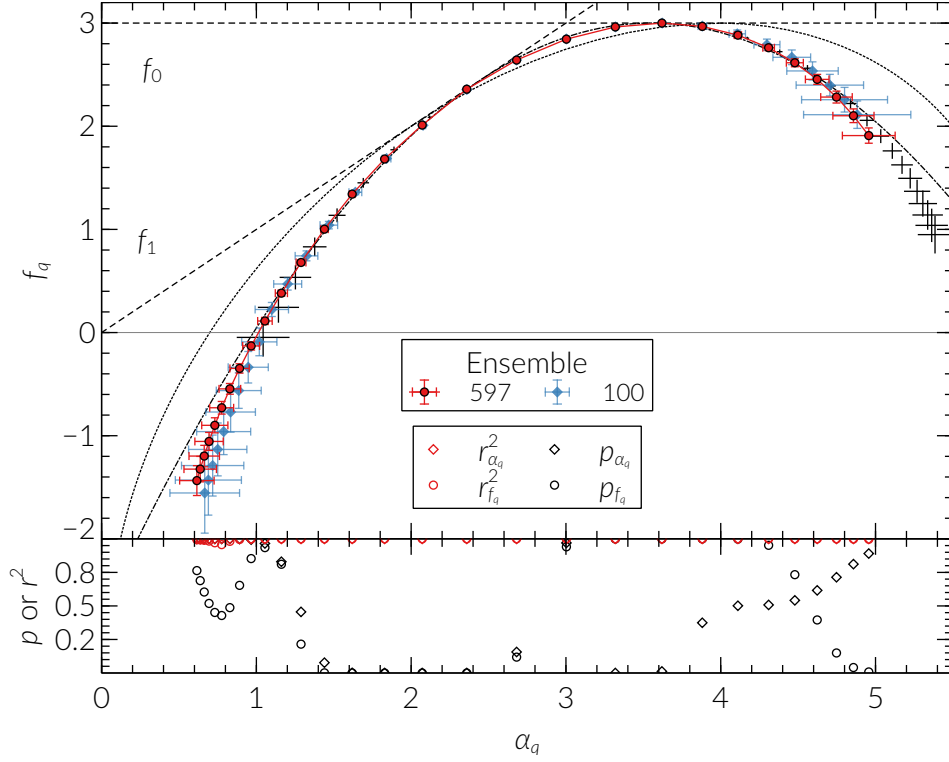


Figure 4.2: Above: singularity spectrum for $L = 22$ and $N_S = 140$, sampled for values of q from -2 to 5 in steps of $1/4$ (increasing from right to left) at energy -0.249 eV. Blue diamonds show the results for the ensemble of the first 100 disorder realisations, while red circles indicate the results from all available realisations (597). Simple error bars, without data point, indicate the symmetrised spectrum to the full ensemble. Dashed lines indicate the functions $f_0 \equiv D$ and $f_1(\alpha) = \alpha$. The dotted line indicates the spectrum for the Anderson model at criticality, reproduced from Rodriguez et al. (2008), while the dot-dashed line indicates the fit to the parabolic approximation (4.18). Below: linear correlation coefficient r^2 (red) and quality of fit p (black) for the linear fits used to extrapolate the thermodynamic limit of α_q (diamonds) and f_q (circles), shown in Fig. 4.3.




Figure 4.3: Linear fits used to produce the data plotted in Fig. 4.2. The slopes of the lines yield α_q (panel above) and f_q (below). For clarity we only show data for integer values of q from -2 (red) to 5 (grey), with data for $q = 0$ highlighted with a full symbol.




Figure 4.4: Singularity spectrum for $L = 22$ and $N_S = 230$, sampled for values of q from -2 to 5 in steps of $1/4$ (increasing from right to left). The ensemble contains 500 realisations. Green circles indicate the average over the ensemble of states near the Fermi energy $\epsilon_F = 0$, while red squares indicate the ensemble over states closest to -0.320 eV. The corresponding symmetrised spectra are indicated with the same colours by the error bars only.

Chapter 5

Finite-size scaling of the transition

5.1 Scaling theory of localisation

The “scaling theory of localisation” was formulated by Abrahams et al. (1979) and is based on the ideas developed in the 1970s, most notably by Landauer (1970), Thouless (1974) and Wegner (1976). In Ch. 4 we saw that, at the Anderson transition, the correlation length ξ diverges when approaching the critical point and the state of the system shows self-similarity, at least for a finite range of coarse grainings. Following renormalisation group arguments (Wilson, 1971), Wegner obtains the scaling of the correlation (localisation) length as

$$\xi \sim |w|^{-\nu}. \quad (5.1)$$

While we refer to Wegner (1976) for the formal proof, an intuitive argument to derive this relation proceeds as follows. Let us consider a transformation that rescales lengths r and energies ε such that the density of states (per energy and volume) remains constant: $r \sim b^{-1}$ and $\varepsilon \sim b^d$, where b is the resolution and d is the dimensionality of the space (cf. Ch. 4 for the notation). Let w indicate a (perturbatively small) distance from the critical point, with $w > 0$ describing the delocalised phase, and let us assume that w also rescales with the resolution b with an exponent y such that $w \sim b^y$. This scaling relation can be inverted (i.e. we look at how w changes to establish at which resolution we are studying the transition) to write $b \sim |w|^{1/y}$, such that the scaling relation for lengths (including ξ) can be rewritten as $r \sim |w|^{-1/y} = |w|^{-\nu}$, if we set $y = 1/\nu$.

In the same work, Wegner (1976) uses the Kubo-Greenwood formula (Kubo,

1957; Greenwood, 1958) to derive the scaling of the conductivity as

$$\sigma \sim w^{(d-2)\nu} \quad w > 0. \quad (5.2)$$

The constraint $w > 0$ ensures that the conductivity is non-vanishing only in the extended regime. The scaling exponent of the conductivity is often defined as $s = (d - 2)\nu$, which constitutes a *hyperscaling relation* because it involves the dimensionality d of the system (Cardy, 1996; Herbut, 2007). Hyperscaling holds within a finite critical dimension d_c , but for the metal-insulator transition it holds in any dimension (Imada et al., 1998). Notice that for $d = 3$, we have $s = \nu$.

In the late 1970s, it became apparent that a scaling theory of localisation should be based on the *conductance* G , an extensive property, rather than the conductivity, which is intensive, i.e.

$$G = \sigma L^{d-2}. \quad (5.3)$$

In their scaling argument for G , Thouless and co-workers (Edwards and Thouless, 1972; Thouless, 1974) consider a system of size L^d and replicate it to construct a system of size $(2L)^d$, so that the eigenstate of the larger sample is a linear combination of the eigenstates of the smaller sample. They then compute the energy difference ΔE obtained by changing the boundary conditions from periodic to antiperiodic, with the idea that a state localised in one of the subsystem is barely affected by boundary effects. Finally, they compare ΔE to the average spacing δW between energy levels and show that $\Delta E/\delta W$ becomes exponentially small in the localised phase, while it is large for extended states. This “Thouless number” $\Delta E/\delta W$ is actually the *dimensionless* conductance:

$$g = \frac{\Delta E}{\delta W} = \frac{G}{e^2/h} \propto L^{d-2}. \quad (5.4)$$

In their pioneering work, Abrahams et al. (1979) assume that g is the only relevant variable in a *scaling theory of localisation*, and that it obeys the scaling relation

$$g(bL) = f(b, g(L)) \quad \text{or, equivalently,} \quad \beta[g(L)] = \frac{d \log g(L)}{d \log L}. \quad (5.5)$$

Here β (or, equivalently, f) is an analytic function (at least near the transition) that depends explicitly only on the dimensionless conductance, and *not* on the system size, energy or microscopic details like the disorder realisation.




Figure 5.1: Sketch of the asymptotic behaviour of the β function of Eq. (5.5) for dimensions $d = 1$ to 3. In the $g \rightarrow +\infty$ limit we have $\beta \rightarrow d - 2$, while for $g \rightarrow -\infty$ we have $\beta \propto \log g$. The black dot indicates the critical conductance g_c such that $\beta(g_c) = 0$.

While a functional form of β is unknown, we can have a qualitative idea of its behaviour by studying some asymptotic cases, which we show in Fig. 5.1. For $\beta > 0$, the logarithmic derivative of $g(L)$ is positive, namely the dimensionless conductance increases with system size. This is the metallic behaviour of Eq. (5.3), which means that, for large g , $\beta(g) = d - 2$. If $\beta < 0$, instead, g decreases with system size. This is the behaviour in the localised phase (cf. Eq. 3.6), where the Thouless number becomes exponentially small: $g(L) \propto \exp(-L/\xi)$, which yields $\beta(g) \propto \log g$. For $\beta = 0$, the dimensionless conductance is invariant with system size: this identifies the critical point of the transition. One of the most acclaimed results of the scaling theory of localisation is that it predicts a phase transition only for $d > 2$, where, at g_c , β crosses from negative to positive values (see Ch. 1). Since β is always negative for $d \leq 2$, at least assuming monotonicity, the eigenstates of an “Anderson system” ($T = 0$, non-interacting electrons, without magnetic scattering) are always localised.

5.2 The problem with finite systems

The *critical phenomena* we have discussed in Ch. 4 and in the previous section, starting from the divergence of the correlation length (5.1), are formally defined only in the *thermodynamic limit* (Van Hove, 1949), namely when the system size

is infinite¹. In *finite* systems, instead, the phase transition loses the sharpness stemming from the non-analyticity, since the partition function of the system, as a finite sum of analytic functions, is itself analytic (Herbut, 2007). Of course, real macroscopic systems have such a large number of constituents, often compared to the Avogadro number $N_A \approx 6.022 \times 10^{23} \text{ mol}^{-1}$, that they are effectively *very close* to the thermodynamic limit. As of 2018, no computer simulation is able to get even close to simulating physical systems of macroscopic size. In this work, specifically, we study systems with 10^3 – 10^4 atoms.

Let us, for simplicity, consider a cubic system of linear size L and measure on it an observable a that assumes the value a_∞ in the thermodynamic limit. In our finite system we expect to see $a_L = a_\infty + \delta a_L$, with δa_L a finite-size perturbation vanishing in the limit $L \rightarrow \infty$ and containing information on the geometry and boundary conditions (Brankov et al., 2000). This correction becomes relevant in a second-order phase transition, where properties of the system are determined by the correlation between the microscopic degrees of freedom and the correlation length ξ cannot grow beyond the system size L .

The procedure that relates a_L to its thermodynamic limit a_∞ is called *finite-size scaling* (FSS) analysis, and was established by Fisher and Barber (1972). For the analysis that follows we rely on the assumptions of *phenomenological* FSS: L/ξ is the only relevant variable necessary to study finite-size effects, and these are appreciable, as described in the previous paragraph, when $L \sim \xi$ (Brankov et al., 2000). The FSS approach was used by Pichard and Sarma (1981) and MacKinnon and Kramer (1981) to test the one-parameter scaling model of Eq. (5.5) for the Anderson model. At a given dimension d , the (finite-size) correlation length at disorder w and length L is given by

$$\xi_L(w) = \xi_\infty(w) f_d \left(\frac{L}{\xi_\infty(w)} \right), \quad (5.6)$$

where ξ_∞ is the correlation length in the thermodynamic limit and the scaling function f_d is different for each dimension. By changing variables from L/ξ_∞ to g , it is possible to relate f_d to β (MacKinnon and Kramer, 1981). In fact, assuming the scaling theory of localisation (5.5) is equivalent to assuming the scaling hypothesis (5.6) for the critical behaviour, given the existence of a finite conductance g_c at

¹More precisely, in the thermodynamic limit the number of atoms N and the volume of a system $V = L^3$ (assumed cubic for simplicity) diverge with constant concentration $\rho = N/V$. Since this is realised in periodic structures like crystals, we will often just mention $L \rightarrow \infty$ as the thermodynamic limit, as is often found in the literature.

the transition (Dobrosavljević et al., 2012).

5.3 Finite-size scaling of multifractal exponents

More generally we assume that the scaling behaviour of a quantity Γ around the critical point in a system of size L can be described by an analytic function $f_\Gamma(L/\xi)$. If we identify the source of disorder in the (rescaled) concentration of dopants $w = (n - n_c)/n_c$, from Eq. (5.1) we can write

$$L/\xi(w) \sim L|w|^\nu. \quad (5.7)$$

It is customary to rewrite f_Γ as a function of $wL^{1/\nu}$, where, using the vocabulary of the renormalisation group approach (Cardy, 1996), w is the *relevant field* and $1/\nu$ is the associated *relevant exponent*.

The characterisation of the transition we present in this work is based on the statistical properties of the electronic wave functions. The candidate quantities for Γ , hence, are the multifractal exponents (4.15), in particular α_q . We now momentarily forget to question whether such quantities are suitable for a scaling analysis, and derive expressions for potential scaling functions f_Γ .

5.3.1 Derivation of scaling functions

To extend the multifractal analysis beyond the critical point we follow the work by Rodriguez et al. (2011), which is itself based on Yakubo and Ono (1998). The starting assumption is that the partition sum R_q of (4.6) for the ensemble of realisations at fixed w , L and λ scales like (5.6):

$$\langle R_q \rangle(w, L, \lambda) = \lambda^{\tau_q^\infty} \mathcal{R}_q(L/\xi(w), \lambda), \quad (5.8)$$

where τ_q^∞ is the mass exponent τ_q in the thermodynamic limit. Compared to (4.13), the scaling function \mathcal{R}_q includes the deviation from the critical point ($w \neq 0$), the finite-size effects (finite L) and coarse-graining effects ($\lambda > 0$). The mass exponents are derived from R_q like in Sec. 4.2:

$$\tau_q(w, L, \lambda) = \tau_q^\infty + \frac{q(q-1)}{\log \lambda} \mathcal{T}_q(L/\xi(w), \lambda), \quad (5.9)$$

where \mathcal{T}_q is related to \mathcal{R}_q , and the factor $q(q-1)$ ensures that $\tau_0(w, L, \lambda) = -d$ and $\tau_1(w, L, \lambda) = 0$, $\forall w, L, \lambda$. Analogously, since $\Delta_q = \tau_q - d(q-1)$, we have

$$\Delta_q(w, L, \lambda) = \Delta_q^\infty + \frac{q(q-1)}{\log \lambda} \mathcal{T}_q(L/\xi(w), \lambda), \quad (5.10)$$

while

$$\alpha_q(w, L, \lambda) = \frac{d \tau_q(w, L, \lambda)}{d q} = \alpha_q^\infty + \frac{1}{\log \lambda} \mathcal{A}_q(L/\xi(w), \lambda). \quad (5.11)$$

When we approach the critical point, $w \rightarrow 0$ and $\xi \rightarrow \infty$, the functions \mathcal{T}_q and \mathcal{A}_q attain a finite value such that, e.g. for α_q ,

$$\alpha_q^\lambda = \lim_{w \rightarrow 0} \alpha_q(w, L, \lambda) = \alpha_q^\infty + \frac{\mathcal{A}_q(0, \lambda)}{\log \lambda}, \quad (5.12)$$

where the second term is system-size invariant (depends only on λ) and vanishes in the thermodynamic limit $\lambda \rightarrow 0$.

5.3.2 Validity of the scaling assumption

The question that arises when dealing with finite systems is whether the wave function is still a multifractal. Formally speaking, the wave function is a true multifractal only at the critical point (Sec. 4.1). For a finite L , however, the effective critical point shifts away as $L^{-1/\nu}$ from its thermodynamic limit (Cardy, 1996). Luckily this is not a problem, since, as shown by Cuevas and Kravtsov (2007), states on the two sides of the transition still show multifractal features characteristic of a critical wave function.

Now that we can construct a multifractal measure from the wave functions away from the critical point, we can actually check the most important assumption we have taken so far, namely that a localisation-delocalisation transition occurs in our model. From the discussion in Ch. 4, in particular Eq. (4.8), we know that the histogram distribution $N_\lambda(\alpha)$ of the measure α depends, at the critical point, only on the coarse-graining $\lambda = l/L$, rather than separately on the system size L and the box size l . At the critical point, then, $N_\lambda(\alpha)$ has the same shape for any L , provided that the wave functions are coarse-grained with the matching l box size. The dependence of $N_\lambda(\alpha)$ on L gradually reappears away from the critical point, where in the strong (weak) disorder regime, larger systems become more localized (delocalized). This is shown in Fig. 5.2, where we plot the (ensemble) probability




Figure 5.2: Ensemble PDF of the multifractal measure α at coarse-graining $\lambda = 1/2$ and energy $\varepsilon - \varepsilon_F = -0.249$ eV, as a function of the concentration n (in units of 10^{20} cm^{-3}), for two system sizes $L^3 = 4096$ (blue dots) and 10648 (red crosses). For clarity we show the histogram for three concentrations: before the transition ($n = 4.6 \times 10^{20} \text{ cm}^{-3}$), near the critical point ($6.8 \times 10^{20} \text{ cm}^{-3}$), and after ($8.8 \times 10^{20} \text{ cm}^{-3}$). The critical point ($n_c = 6.7 \times 10^{20} \text{ cm}^{-3}$) is indicated by a black dashed line and the value used is from table A.3. On the bottom plane we show the position of the average α_0 also for the intermediate concentrations, connected by lines to guide the eye. We use again blue dots with a solid line for $L^3 = 4096$ and red crosses with a dashed line for 10648 . Reproduced from the Supplemental Materials to Carnio et al. (2017).

distribution function (normalised histogram) $\text{PDF}(\alpha) = N_\lambda(\alpha)\lambda^d/\mathcal{N}$ at $\lambda = 1/2$ for three values of n : the lowest in the localized regime, the intermediate close to the critical point and the highest in the delocalised regime. The ensemble $\text{PDF}(\alpha)$ is built by filling one single histogram with the \mathcal{N} available wave functions.

Because of the limited spread in system sizes and the large common coarse-graining, the difference in the PDFs outside the critical point is not very well pronounced in Fig. 5.2. An alternative check consists in fixing the system size and study how the PDF renormalises with the coarse graining (Lindinger and Rodríguez, 2017). As presented in Sec. 5.1, the “MacKinnon-Kramer” variable ξ/L (MacKinnon and Kramer, 1981) scales like λ^{-1} , which implies that, with increasing λ , ξ/L becomes smaller. Physically this means that, upon coarse-graining, localised (delocalised) states become more localised (delocalised), or, equivalently, that the renormalisation flow rescales the disorder away from its critical value,² if a phase transition, and hence a critical point, exists. We verify this in Fig. 5.3: upon increasing the box size l in a system of $L^3 = 4096$ atoms, the PDF’s move in opposite directions.

5.4 Scaling at fixed coarse-graining ratio

We now present how we fit the data from our simulations (Sec. 4.4) to scaling functions from the previous section. In particular, the function in Eq. (5.11) contemplates scaling in two variables, L and λ (or, equivalently, $l = \lambda L$). A possible simplification involves considering data at a fixed λ , which is then expected to scale only with system size (Rodríguez et al., 2011).

The coarse-graining procedure is straightforward: at resolution λ , we define λ^{-3} boxes with volume $(\lambda L)^3$ and calculate the corresponding measures μ_i , as described in Sec. 4.2, from the sites of the wave function that fall into the each box. This implies that, if $l = \lambda L$ is not an integer, the number of points from the original wave function that falls in each box will be different. This procedure differs from that discussed in Sec. 4.3, where each box contains the same number of points from the original wave function. With this latter partitioning method we cannot obtain a common value of λ between the system sizes.

In the following, only $\lambda = 1/2$ yields an equal number of points in each box. In the other cases, what might cause a problem is the accumulation of probability in a certain box due to the larger number of points, so that the coarse grained wave

²The value of disorder is rescaled to lower (higher) values in the localised (delocalised) phase.




Figure 5.3: Ensemble PDF of the multifractal measure α at energy $\varepsilon - \varepsilon_F = -0.249$ eV for $L^3 = 4096$ atoms. In blue we show the PDF for $N_S = 20$, in orange for $N_S = 100$. Filled symbols indicate a coarse-graining with box size $l = 1$, while empty symbols connected by a dotted line indicate $l = 4$.

function looks more localised in said box. We expect this type of fluctuation to be compensated in large ensembles, where the chances of finding a low-probability box increases.

5.4.1 Fits and scaling

At a fixed λ , the only variable left in Eq. (5.11) is L/ξ , and hence we can work with the equivalent function $\tilde{\mathcal{A}}_q(\rho L^{1/\nu})$ to be determined. Notice that, compared to Sec. 5.3, our scaling function depends on a new relevant variable ρ that takes into account possible non-linear dependencies on w :

$$\rho(w) = w + \sum_{i=2}^{m_\rho} b_i w^i. \quad (5.13)$$

Assuming its analyticity, the scaling function is also Taylor-expanded around $\rho L^{1/\nu} = 0$, i.e. at the critical point $\rho = 0$:

$$\tilde{\mathcal{A}}_q(\rho L^{1/\nu}) = \sum_{i=0}^{n_L} a_i \rho^i L^{i/\nu}. \quad (5.14)$$

The zeroth-order term a_0 , in the case of fixed λ , is equal to α_q^λ of Eq. (5.12). Like in MacKinnon and Kramer (1981), the function $\tilde{\mathcal{A}}_q$ can be used to collapse the data on two branches given by

$$\tilde{\mathcal{A}}_q(\pm(L/\xi)^{1/\nu}), \quad (5.15)$$

where the positive (negative) sign indicates, in our model, the delocalised (localised) phase (see the example of Fig. 5.4).

Following Rodriguez et al. (2011) we fit the data $\alpha_q(w, L, \lambda)$ against the scaling function (5.14), for increasing expansion orders n_L and m_ρ , to determine the values of the parameters $\{a_i\}$, $\{b_i\}$, n_c (implicit in w) and ν , for a total of $N_p = n_L + m_\rho + 2$ parameters. We fit the data using the method of least squares minimization, i.e. we minimize the quantity

$$\chi^2 = \sum_{i=1}^{N_D} \frac{[y_i - f(x_i)]^2}{\sigma_i^2}, \quad (5.16)$$

where y_i are the N_D data points corresponding to the x_i initial conditions, σ_i their corresponding uncertainty and f is the fitting function we are considering (Press et al., 2007). We also calculate the goodness-of-fit p as

$$p = \frac{\Gamma\left(\frac{1}{2}(N_p - N_D), \frac{1}{2}\chi^2\right)}{\Gamma\left(\frac{1}{2}(N_p - N_D)\right)}, \quad (5.17)$$

where $\Gamma(x)$ is the Euler Γ function and $\Gamma(a, x)$ is the upper incomplete Γ function (Abramowitz and Stegun, 1970). In Appendix A, and in the figures in Sec. 5.6, we present fits with a goodness-of-fit value $p \geq 0.05$, the value conventionally used in statistical hypothesis testing (Owen and Jones, 1994). For each energy value, we use the smallest concentration interval that yields the smallest uncertainties in n_c and ν . We also make sure that the estimates of the critical parameters do not change, within error bars, with larger concentration intervals (*robust* fits) and when increasing the expansion order of the scaling function (*stable* fits).

5.4.2 An example

As an example, in Fig. 5.4 we show a fit for α_0 at energy -0.249 eV and $\lambda = 1/2$, including $N_D = 106$ data points. From the p value we accept a fit with $n_L = 3$

Table 5.1: Test of the robustness of the fit described in Sec. 5.4.2, containing the description of the symbols. All concentrations are expressed in units of 10^{20} cm^{-3} . The centre of the concentration interval is $n_0 = 6.7 \times 10^{20} \text{ cm}^{-3}$. The star indicates the candidate fit whose robustness is tested.

	n_L	m_ρ	N_P	χ^2	p	n_c	ν	ω
	3	1	6	67	0.55	6.87(30)	1.33(56)	0.15
	3	1	6	86	0.47	6.74(26)	1.39(43)	0.20
★	3	1	6	100	0.48	6.72(23)	1.25(30)	0.25
	3	1	6	120	0.35	6.78(24)	1.44(28)	0.30
	3	1	6	161	0.02	6.77(24)	1.50(26)	0.35

and $m_\rho = 1$ (hence $N_P = 6$), yielding $\chi^2 \approx 100$ and $p \approx 0.48$. The results for the critical parameters are

$$n_c = 6.72(23) \times 10^{20} \text{ cm}^{-3} \quad \nu = 1.25(30). \quad (5.18)$$

In table 5.1 we test the robustness of the fits, i.e. how the critical parameters change with the width of the critical region, measured as

$$\omega = \frac{n_{\max} - n_{\min}}{n_{\max} + n_{\min}}. \quad (5.19)$$

The parameters from all fits are compatible within the error bars. As a rule of thumb, we choose the smallest spread that minimizes the uncertainties. In this example, for $\omega < 0.25$ we obtain significantly higher uncertainties, while for $\omega > 0.25$ the p value decreases lower than the acceptance threshold of 0.05. From this pool, therefore, we choose $\omega = 0.25$ as the best fit. This fit is then tested for stability, namely we increase n_L and m_ρ , separately and together, as reported in table 5.2. This second set of fits shows similar p values and compatible critical parameters, hence we accept the candidate fit as robust and stable.

In Fig. 5.4 we show the data and the scaling function (5.14) for the four different values of L . By definition the curves meet at n_c , where the system shows size invariance. From the fitted values of n_c and ν we replot the data as a function of L/ξ (5.7) and show that it collapses on the curves $\tilde{\mathcal{A}}_0(\pm(L/\xi)^{1/\nu})$.

Table 5.2: Test of the stability of the fit described in Sec. 5.4.2 and highlighted in Table 5.1. All concentrations are expressed in units of 10^{20} cm^{-3} . The centre of the concentration interval is $n_0 = 6.7 \times 10^{20} \text{ cm}^{-3}$ and its width is $\omega = 0.25$. The star indicates the candidate fit whose stability is tested.

	n_L	m_ρ	N_P	χ^2	p	n_c	ν
★	3	1	6	100	0.48	6.72(23)	1.25(30)
	4	1	7	100	0.46	6.71(23)	1.26(30)
	3	2	7	99	0.48	6.79(25)	1.19(25)
	4	2	8	99	0.47	6.82(25)	1.18(25)

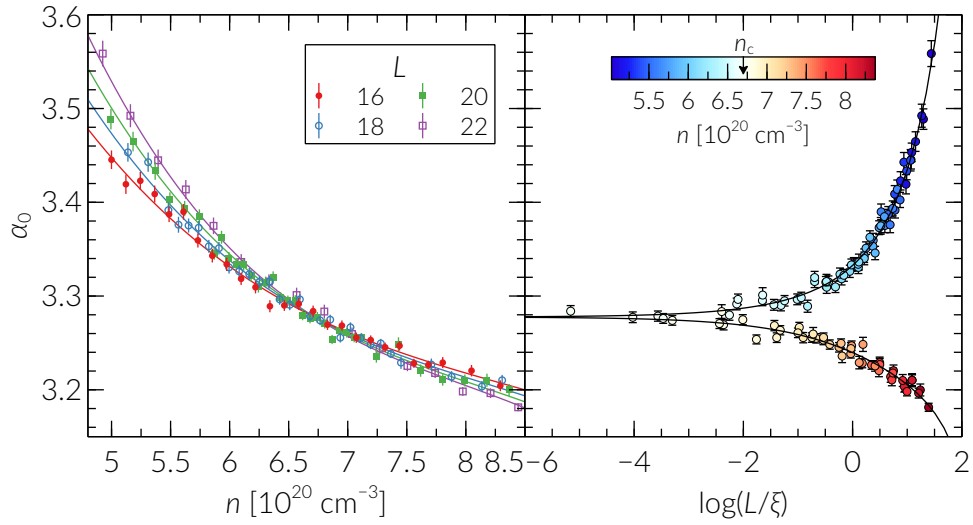


Figure 5.4: Scaling of α_0 for energy -0.249 eV and $\lambda = 1/2$, with points indicating the numeric values obtained from the ensemble average. Left: fitting of the data points α_0 for different concentration n and system size L , indicated by colour. The same colours are used to plot the scaling functions $\tilde{\mathcal{A}}_0(\rho L^{1/\nu})$ for the different values of L . Right: same data points with abscissa rescaled as L/ξ (in log scale) and colours indicating the value of n for each data point. The critical concentration is indicated as n_c on the scale. The underlying fit is given by $\tilde{\mathcal{A}}_0(\pm(L/\xi)^{1/\nu})$. Reproduced from the Supplemental Materials to Carnio et al. (2017).

5.5 Other scaling possibilities

5.5.1 Inclusion of irrelevant corrections

If the data series do not meet in the single point of Eq. (5.12), we can consider a scaling function that depends additionally on an *irrelevant* field η with a corresponding irrelevant exponent $y < 0$, i.e. $\tilde{\mathcal{A}}_q(\rho L^{1/\nu}, \eta L^y)$. While we could fit a two-dimensional Taylor expansion in the two variables as done in Slevin and Ohtsuki (1999), we follow the approach in Rodriguez et al. (2011) of expanding the scaling function in the irrelevant variable and truncating the expansion after the first order:

$$\tilde{\mathcal{A}}_q(\rho L^{1/\nu}, \eta L^y) = \tilde{\mathcal{A}}_q^0(\rho L^{1/\nu}) + \eta L^y \tilde{\mathcal{A}}_q^1(\rho L^{1/\nu}). \quad (5.20)$$

Like in (5.14), the one-parameter scaling functions $\tilde{\mathcal{A}}_q^0$ and $\tilde{\mathcal{A}}_q^1$ are expanded in their arguments up to orders n_0 and n_1 , respectively. Similarly, the irrelevant field is expanded as

$$\eta(w) = 1 + \sum_{i=1}^{m_\eta} c_i w^i. \quad (5.21)$$

The zeroth-order term is just a factor in front of the expansion parameters of $\tilde{\mathcal{A}}_q^1$, hence it has been set to 1 to minimize the number of parameters. This amounts, then, to $N_p = n_0 + n_1 + m_\rho + m_\eta + 4$. Because of its accuracy in estimating the critical parameters of the transition, this method has also been used to study the metal-insulator transition in other Anderson systems (Puschmann et al., 2015; Moore et al., 2017).

Let us clarify here that, for most energy channels at $\lambda = 1/2$ and $1/4$ we do *not* see a shift in the crossing point of the data series. In the following we try fitting the function (5.20) to our data to *exclude* the need for an irrelevant term. As we will repeat again in the following, for large λ and small system sizes, fitting with (5.14) is adequate and yields the usual estimate $\nu \approx 1.6$ in studies on the Anderson localisation in different universality classes (Slevin and Ohtsuki, 2014; Devakul and Huse, 2017).

As shown in table 5.3, fits using (5.20) converge and have $p > 0.05$, but do not consistently meet the stability criterion for the available system sizes. For some of them y is very large in magnitude and has a vanishing uncertainty. In fact, in the limit $y \rightarrow -\infty$ the irrelevant contribution in (5.20) disappears and we are effectively fitting the data to (5.14), but with unnecessary additional parameters.

This is indeed the case of the fit indicated with a star, which corresponds to the best fit obtained in table 5.1. In other cases, instead, the uncertainty on y (or ν) is of the same order of magnitude, or bigger, than the value itself.

As expected, since with our data we cannot see a shift in the critical point, it is not possible to reliably estimate an irrelevant exponent. This is consistent with the renormalisation group argument (Cardy, 1996) that irrelevant variables vanish with large enough coarse-graining. This effect is shown also in Lindinger and Rodríguez (2017): a comparison between different values of λ shows that, for $\lambda = 1/2$, the irrelevant shift vanishes and the crossing point of the data series approaches the critical point. In the same work, the values of the critical disorder and the critical exponent remain consistent for the different coarse grainings, while, as expected, y decreases to more negative values.

For smaller coarse-grainings like $\lambda = 1/4$ and $\lambda = 1/6$, however, we are still unable to estimate the irrelevant exponent, with results undistinguishable from those in table 5.3. In Fig. 5.5 we show the corresponding data sets: the irrelevant shift, if present, cannot be distinguished and, as discussed above, the crossing point shifts to higher concentrations compared to that for $\lambda = 1/2$ (dashed line). This does not necessarily mean that there is no irrelevant shift: as shown in Fig. 5.6, the quality of fit dramatically decreases with $\lambda = 1/6$, meaning either that (5.14) is insufficient to describe the trend in the data or that the error bars are underestimated compared to the intrinsic fluctuations in the ensemble average. In this case, $p \geq 0.05$ is a sufficient criterion to filter out inadequate fits.

The irrelevant shift is probably hidden by two concurrent factors: small system sizes and too few realisations. While the shift should be more visible for small systems (the irrelevant term scales like $L^{-|y|}$), it might still be smaller than the sampling of the concentrations, which cannot be finer than substituting one additional site of the lattice and hence decays like L^{-3} , see (3.11). This also implies that, to have enough data points for a convergent fit, we need to consider a larger concentration range, at least compared to the disorder values used in Rodríguez et al. (2011). As for the number of realisations, our data shows fluctuations larger than the uncertainties associated to the points, which are standard errors of the mean and decay as $\sqrt{\mathcal{N}}$, where \mathcal{N} is the number of realisations. Our ensembles reach a maximum of 1000 realisations, which is an order of magnitude smaller than in the high-precision studies of the non-interacting Anderson transition (Rodríguez et al., 2011; Lindinger and Rodríguez, 2017) due to the additional numerical requirements. This effect is partially mitigated by renormalizing the

Table 5.3: Results for the fits of the data in Fig. 5.4 ($\lambda = 1/2$) with an irrelevant term.

	n_0	n_1	m_ρ	m_η	N_P	χ^2	p	n_c	ν	$-y$
★	3	0	1	0	8	97	0.39	6.73 ± 0.24	1.29 ± 0.35	4.36×10^5
	3	0	2	0	9	94	0.45	7.47 ± 7.22	1.34 ± 1.00	1.69 ± 26.27
	3	1	1	0	9	97	0.36	6.78 ± 0.24	1.29 ± 0.35	5.66×10^5
	3	0	2	0	9	94	0.44	6.83 ± 0.24	1.25 ± 0.29	5.18×10^6
	3	2	1	0	10	96	0.37	7.12 ± 3.31	1.05 ± 1.52	2.98 ± 24.43
	3	0	2	1	10	94	0.41	6.83 ± 0.24	1.25 ± 0.29	3.64×10^6
	3	1	2	0	10	94	0.43	6.93 ± 0.75	1.06 ± 0.92	7.50 ± 51.60
	3	1	1	1	10	96	0.37	6.93 ± 1.68	0.97 ± 1.25	3.03 ± 24.13
	3	1	2	1	11	94	0.38	6.83 ± 0.24	1.25 ± 0.29	277.21
	3	2	1	1	11	97	0.31	6.73 ± 0.24	1.29 ± 0.35	103.75
	3	1	2	1	11	93	0.41	6.99 ± 1.35	0.89 ± 1.93	3.45 ± 34.25
	3	2	1	1	11	94	0.39	7.14 ± 0.94	0.42 ± 0.50	2.05 ± 2.11
	3	2	2	1	12	92	0.43	7.22 ± 0.76	0.37 ± 0.44	1.95 ± 2.12

wave functions: for $\lambda = 1/2$ the renormalised system is defined on $\lambda^{-3} = 8$ sites and the distribution of the GMFE's (4.15) will be broader, as the chance of finding large box-probabilities increases (see, e.g. Fig. 5.3). This, in turn, translates into larger error bars on the values of α_q . The downside, however, is that the estimates of the critical parameters will also have a larger uncertainty.

To conclude, the best analysis we can perform on the data is by fitting the $\lambda = 1/2$ and $\lambda = 1/4$ coarse-grained data with only a relevant variable. By accepting robust and stable fits with $p \geq 0.05$, this method allows to consistently and reliably find best fit estimates.

5.5.2 Two-parameter scaling

Another possibility is to scale the data with both L and λ using Eq. (5.11), or, equivalently, with $l = \lambda L$. This method was successfully established by Rodriguez et al. (2011) because it has the merit of determining, simultaneously, the critical parameters of the transition and the thermodynamic-limit value of the multifractal exponent under study.

Its disadvantage, instead, is that it relies on correlated data – it uses the same ensemble for different coarse grainings – and hence it requires very high precision, i.e. large ensembles. Moreover, as a correction for non-vanishing λ , Eq. (5.11) and its Taylor expansion hold for small λ . As we have small system sizes compared to Rodriguez et al. (2011), we can only access larger values of λ . This might be a reason why we are not able to obtain an acceptable fit using this method.




Figure 5.5: Data sets for $\lambda = 1/4$ (left) and $\lambda = 1/6$ (right) and energy -0.249 eV, in the same format of Fig. 5.4. The estimated crossing points are $n_0 = 7.8 \times 10^{20} \text{ cm}^{-3}$ (left) and $8.5 \times 10^{20} \text{ cm}^{-3}$ (right), while the width of the critical region (5.19) is, in both cases, $\omega = 0.25$. The dashed line indicates the fit estimate $n_c \approx 6.72 \times 10^{20} \text{ cm}^{-3}$ found for $\lambda = 1/2$.




Figure 5.6: The p values obtained from fitting the data at fixed $\lambda = 1/6$ resolution. The horizontal line indicates the threshold of 0.05 above which fits are accepted and hence tested for stability and robustness.

5.6 Phase diagrams

We report now the results obtained from fitting the scaling behaviour of exponents α_0 and α_1 for $\lambda = 1/2$ and $\lambda = 1/4$, as discussed in Sec. 5.4. We fit the ensemble-averaged data near the energy bins $\{E_j\}$ below the Fermi energy (see Sec. 4.4) and listed in Appendix A. We emphasize that, for each energy, we have taken a single wave function from each sample to avoid inter-sample correlations (Rodriguez et al., 2011).

For each fit we obtain the critical concentration n_c and exponent ν and in the following we plot phase diagrams, i.e. how the critical parameters change within the impurity band (IB). The phase diagram obtained when including all four system sizes is shown in Fig. 5.7.

5.6.1 Critical concentration

The set of critical concentrations as a function of energy constitute the *mobility edge* (see Ch. 1): for higher concentrations the wave functions are, on average and in the thermodynamic limit, delocalised, otherwise they are localised.

In Fig. 5.7 the mobility edge appears flat for energies $-0.02 \text{ eV} \lesssim \varepsilon - \varepsilon_F \leq 0$ and then quickly decreases until it reaches a minimum at around -0.1 eV . For lower energies the mobility edge increases again, with smaller but relatively constant slope. For the intermediate energies $-0.15 \text{ eV} \lesssim \varepsilon - \varepsilon_F \lesssim -0.1 \text{ eV}$, some fit results are missing. In this regime the critical concentration is very low, close to the minimum concentration given by having one impurity in the sample, and the available data points in a given range of concentration are few. In these cases, however, the raw data itself shows the critical point shifting to lower concentrations, which confirms that, at these energies, the critical concentration reaches a minimum.

While these results will be discussed in Ch. 6 in the context of the formation of the IB, we point out here already that, from the shape emerging in Fig. 5.7, the possible tails of the band delocalise later than the centre.

5.6.2 Critical exponent

Close to the Fermi energy, we find values $\nu \sim 0.5$, while, for energies $-0.1 \text{ eV} \lesssim \varepsilon - \varepsilon_F \leq 0$, ν increases up to roughly 1. Deeper in the IB, for $\varepsilon - \varepsilon_F \lesssim -0.15 \text{ eV}$, the results for different λ and q are roughly distributed between 1 and 1.5 (see Sec. 5.6.4). For intermediate energies, $-0.15 \text{ eV} \lesssim \varepsilon - \varepsilon_F \lesssim -0.1 \text{ eV}$ the situation

is unclear, for the same reasons discussed for n_c : if the fit converges, is accepted and shows a critical concentration, the estimate of ν fluctuates and shows large error bars: a reflection of the too few and fluctuating data points that underlie the fit. We might reduce the uncertainties by computing more realisations, or, as we mention also in Ch. 7, by considering larger system sizes. In the latter case, wave functions are sampled on larger domains, which improves the estimation of the multifractal exponents via Eq. (4.13).

The shape of the phase diagram for ν suggests the existence of at least two different regimes: close to the Fermi energy we find the value $\nu \sim 0.5$ found also in the experiments on Si:P and other uncompensated semiconductors discussed in Ch. 1. Moving deeper into the impurity band, instead, this behaviour changes (roughly continuously, at least at the top of the band), until we reach a relatively stable value for lower energies. The phase diagram stops at around -0.33 eV because, in this work, we generate data up to around $2.5 \times 10^{21} \text{ cm}^{-3}$ (Table 4.1). Higher concentrations require more impurities and hence more computation time, as discussed in Sec. 3.5.

5.6.3 Dependence on system size and resolution

We now discuss the stability of the results with respect to the system sizes included in the fitted data, and to the coarse-graining chosen. In all cases, the scaling function (5.14) supports only a common crossing point (no irrelevant shift), so the critical concentration is usually easily resolved and its estimate has a small uncertainty. The critical exponent, on the other hand, is related to the spread in system sizes and hence is more susceptible to the fluctuations in the data. In general, the accuracy of its estimate depends on the number of system sizes available.

In Fig. 5.8 we plot an earlier phase diagram, computed without the data for $L = 22$. To better show the difference, we pick $q = 0$ and further separate the results for $\lambda = 1/2$ (Fig. 5.9) and $\lambda = 1/4$ (Fig. 5.10). As discussed, the addition of a system size does not yield major changes in the estimates of the critical concentration, except maybe near the Fermi energy. According to their uncertainties, however, the estimates at those energies are still compatible. With respect to the critical exponent, instead, the additional system size helps reducing the uncertainty on the estimate, as discussed above. While the estimates with $L = 22$ are compatible, in the case of $\lambda = 1/4$ (Fig. 5.10) they seem to be consistently slightly lower than those with $L \leq 20$. This seems to be the case also for $\lambda = 1/2$ for the energies at the top of the IB. A bigger difference appears, instead, between the two resolutions.

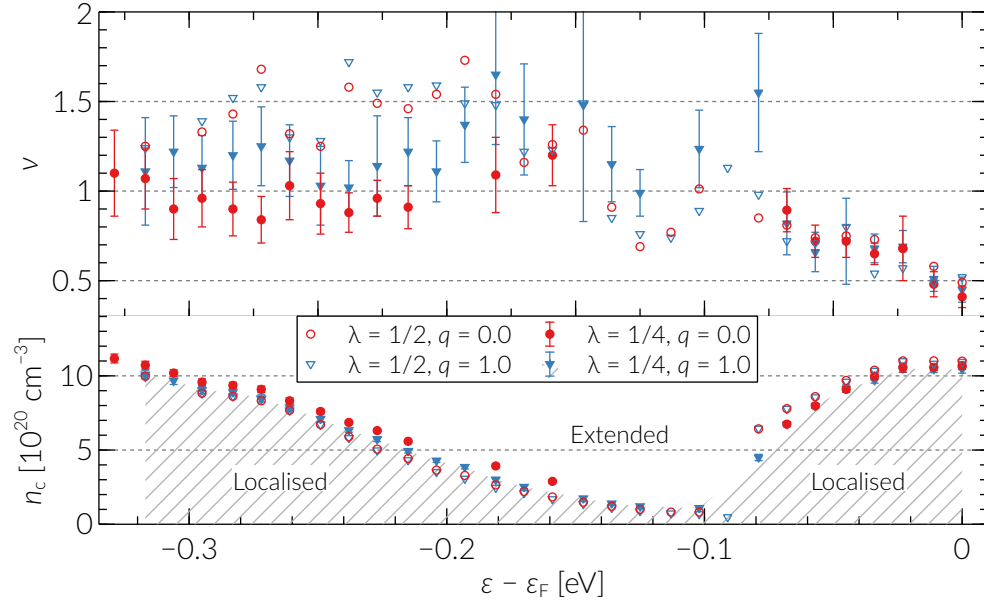


Figure 5.7: Phase diagram with data up to $L = 22$. Critical concentrations n_c and exponents ν as a function of the energy ε from the Fermi level ε_F , for $q = 0$ (red circles) and $q = 1$ (blue triangles). Full and open symbols show, respectively, the results for $\lambda = 1/4$ and $\lambda = 1/2$ coarse-grainings. The error bars, shown only for $\lambda = 1/4$ and if larger than symbol size, represent the 95% confidence level on the fit parameters. The error bars for $\lambda = 1/2$ are of the same order of magnitude as for $\lambda = 1/4$ and are omitted for clarity. The shading indicates the concentrations and energies where the average wave function in the thermodynamic limit is localised. Reproduced from Carnio et al. (2017).

The shift to higher values of n_c when decreasing λ has already been discussed in Sec. 5.5.1 and shown in Fig. 5.5. It is intriguing to see that this indeed happens for $\varepsilon - \varepsilon_F \lesssim -0.1$ eV, while for $\varepsilon - \varepsilon_F \gtrsim -0.1$ eV we observe the opposite behaviour. We have no explanation for this effect.

The changes in the critical exponent, instead, depend on the energy. The increasing trend from 0.5 to 1 at the top of the band, $-0.1 \text{ eV} \lesssim \varepsilon - \varepsilon_F \leq 0$, is present for both coarse-grainings but can be seen more sharply for $\lambda = 1/4$. The “plateau” regime for $\varepsilon - \varepsilon_F \lesssim -0.15$ eV, instead, sees a decrease from approximately 1.5 to 1. Due to the large uncertainties, however, the estimates for λ are compatible. Since the addition of a system size leads to the reduction of the uncertainties, without changing the general trend that we observe, the reduction of ν near the Fermi energy is not likely to be a finite-size effect.




Figure 5.8: Phase diagram with data up to $L = 20$. Symbols, colours, and error bars are used like in Fig. 5.7.




Figure 5.9: Comparison between the fit results when including system sizes up to $L = 20$ (red circles) and $L = 22$ (blue triangles), for $\lambda = 1/2$ and $q = 0$. The error bars represent the 95% confidence level on the fit parameters.




Figure 5.10: Comparison between the fit results when including system sizes up to $L = 20$ (red circles) and $L = 22$ (blue triangles), for $\lambda = 1/4$ and $q = 0$. The error bars represent the 95% confidence level on the fit parameters.

5.6.4 Dependence on multifractal moments

In Fig. 5.11 we compare the phase diagram for $\lambda = 1/2$ and three different values of q . The critical concentration for $q = -1$ essentially follows the results obtained for $q = 0$ and $q = 1$, with only slightly higher values in the middle of the band and a better resolution near $\varepsilon - \varepsilon_F \approx -0.1$ eV. For the initial transient at -0.1 eV $\lesssim \varepsilon - \varepsilon_F \leq 0$, the estimates of the critical exponent are in good agreement for all values of q . At lower energies, instead, the estimates for $q = -1$ are significantly lower, but still agree in most cases. It is also not clear whether they have an increasing trend or they fall on a plateau as seems to happen for $q = 0$ and 1. Moreover, from this comparison seems to emerge a cusp-trend for ν near -0.1 eV. In principle, we do not expect to see any discrepancy in the results for different q values, except for larger estimate uncertainties when studying $q > 1$, where the error bars on the multifractal exponents increase (see Sec. 4.5). From the data we cannot tell whether the difference between q values has a physical explanation or is a statistical artefact due to the small ensemble size.

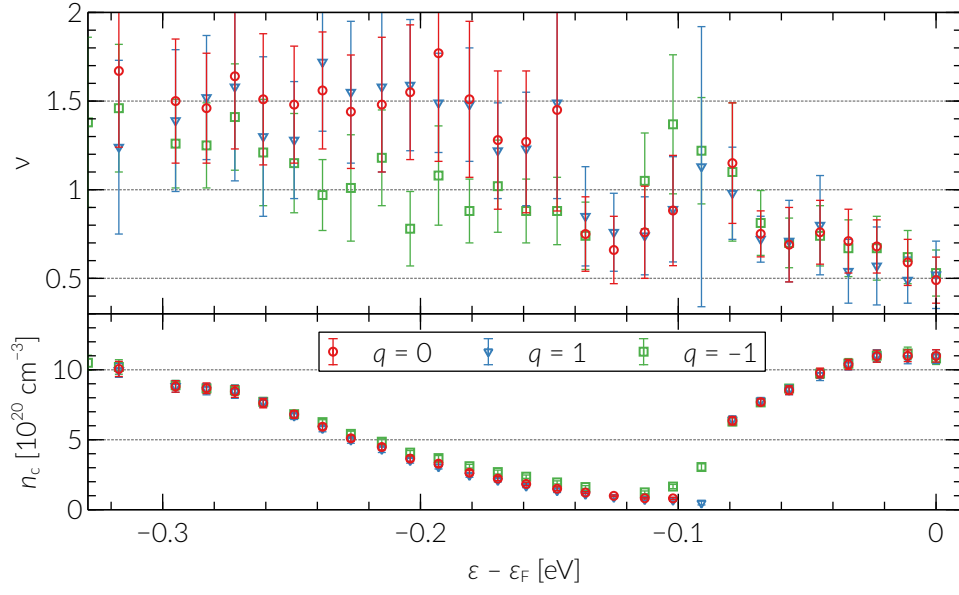


Figure 5.11: Comparison between the fit results when including system sizes up to $L = 22$, for $\lambda = 1/2$ and $q = 0$ (red circles), 1 (blue triangles) and -1 (green squares). The error bars represent the 95% confidence level on the fit parameters.

Chapter 6

Metal-insulator transition in the impurity band

In Ch. 5 we have shown that an Anderson localisation-delocalisation transition occurs in the impurity band. In this Chapter we discuss whether this transition is the backbone of a more general metal-insulator transition in the material. The study of metal-insulator transitions is essentially the work of Sir Neville Mott, hence in this overview we follow his review (Mott, 1990).

6.1 Impurity conduction

Conduction in a doped semiconductor is the result of two competing mechanisms: usual charge transport by conduction-band electrons and donor-to-donor tunnelling of electrons entirely in the impurity band, i.e. without activation into the conduction band. This second process, called *impurity conduction*, depends on the overlap of the wave functions on the impurity sites, and dominates at low temperatures, where the occupancy of the conduction band states vanishes. It was first observed by Hung and Gliessman (1950). Fritzsche (1958) identified two further mechanisms of impurity conduction. For low impurity concentration, transport is possible only when the system is *compensated*, i.e. when a *n*-type semiconductor (doped with donors) contains a minor fraction of acceptors. In this case, the acceptors will accept a number of donated electrons, therefore leaving the same number of donors ionized (unoccupied). The electron in an occupied state can then be thermally excited and tunnel to a vacant state, a process called “hopping” (see Ch. 1).

In uncompensated samples, higher temperatures might excite an electron into hopping among occupied states. The activation energy of this process is seen to decrease for increasing concentration and vanishes at the critical concentration n_c . As we know from Ch. 5, at this concentration states become extended and conduction becomes metallic.

For highly-doped uncompensated semiconductors, the experiment by Alexander and Holcomb (1968) shows the existence of both n_c and an additional characteristic concentration n_{CB} , beyond which the Fermi level moves into the conduction band and, at that energy, electrons show conduction-band properties. In all materials considered, $n_{CB} > n_c$, namely, by increasing concentration, the impurity band first delocalises and then merges with the conduction band. These two processes are independent: the wave function at the Fermi energy might be localised around the donors without there being a gap between impurity and conduction band (Alexander and Holcomb, 1968; Mott, 1990).

So far we have not specified which localisation mechanism underlies the metal-insulator transition. In fact, Mott was convinced initially that the localisation transition was driven by the Hubbard energy¹ U , making it a Mott transition. After reviewing the calculations of Bhatt and Rice (1981) and the experimental observations of Hirsch and Holcomb (1987), however, Mott concluded that, at least in materials like silicon and germanium, with or without compensation, an Anderson transition (with long-range interactions) occurs.

In the rest of the Chapter we will follow the same steps to verify whether our *ab-initio*-based model shows a metal-insulator transition in the impurity band. We already know that the localisation-delocalisation transition is of Anderson type: as a single-particle approach, DFT cannot access electronic correlations underlying Mott physics (see Sec. 2.1.3). From the phase diagram in Fig. 5.7, we also suspect the existence of an impurity band, whose density of states we now properly show. What remains to be checked, then, is the existence of a concentration n_{CB} where the impurity and conduction band merge.

6.2 Formation of the impurity band

From the eigenvalues ε_j in Eq. (3.9), rescaled to the Fermi energy ε_F for each realisation (see Sec. 4.4), we can compute the average density of states (DOS) and observe the formation of the impurity band (IB). Let \mathcal{E}_i be the rescaled spectrum

¹The cost of double occupancy of the same lattice site.

for realisation i for a given system size L and number of impurities N_S . The average DOS is calculated by summing over Gaussian distributions centred on each rescaled eigenvalue $\tilde{\varepsilon} \in \mathcal{E} = \bigcup_i \mathcal{E}_i$ with a standard deviation (smearing) σ :

$$\text{DOS}(\varepsilon) = \frac{1}{\mathcal{N} \sqrt{2\pi\sigma^2}} \sum_{\tilde{\varepsilon} \in \mathcal{E}} e^{-\frac{(\varepsilon - \tilde{\varepsilon})^2}{2\sigma^2}}, \quad (6.1)$$

where \mathcal{N} indicates the number of realisations and ensures that the integrated DOS yields the number of impurities. The DOS shown in Fig. 6.1 is calculated by summing over Gaussian distributions of standard deviation $\sigma = 0.05 \text{ mHa} = 1.36 \text{ meV}$. With increasing concentration of impurities, the IB forms as a peak in the band gap and extends rightwards until it merges with the conduction band (CB), and leftwards towards the valence band (VB) in a long tail. In Fig. 6.1 we have also indicated the energy regions where the states become on average delocalised in the thermodynamic limit, based on the results of Ch. 5. The mobility edges move outwards from the centre of the band, asymmetrically in the direction of the long lower-energy tail. This behaviour descends from the asymmetry in the phase diagram of n_c itself (Fig. 5.7).

The characterisation of the DOS is an important problem in the study of the conduction properties of the impurity band. Early studies based on one-electron tight-binding Hamiltonians (Cyrot-Lackmann and Gaspard, 1974) also find an impurity level that broadens, with increasing concentration, into an asymmetric band with a long tail towards the valence band. In a recent study on electron transport in the IB of a doped semiconductor, Wellens and Jalabert (2016) find qualitatively similar results using a self-consistent diagrammatic perturbation approach. To the best of our knowledge, the result shown Fig. 6.1 is the quantitatively most detailed description of the IB of a doped semiconductor.

The IB plays a central role in semiconductor-based solar cells. As shown in Luque and Martí (1997), intermediate energy levels in the band gap, connected to the valence and conduction bands by optical transitions, can yield very high efficiencies in idealised solar cells. Deep-level impurities, like sulphur in silicon, are therefore excellent candidates to capture low-energy photons (Winkler et al., 2011). When the impurity states are localised, however, they act as nonradiative recombination centres, i.e. they promote electron-hole recombination with the emission of a phonon, rather than a photon, as desired in a photovoltaic device. Luque et al. (2006) show that the delocalisation of the impurity state suppresses this unwanted recombination process. The localisation-delocalisation transition




Figure 6.1: DOS of the IB for 4096 atoms at three different concentrations. The shading indicates energies where states are on average delocalised (in the $L \rightarrow \infty$ limit), according to Fig. 5.7. The delocalised CB states, separated by a vertical dashed line at ε_F , are also shaded. Crossed shading indicates states that might be delocalised, but are outside our concentration range. Reproduced from Carnio et al. (2017).

in impurity bands, then, is essential to reach the efficiency computed for idealised solar cells.

6.3 Band gap closing and the metal-insulator transition

In Fig. 6.1 we can see that the impurity (IB) and conduction band (CB) merge when increasing the concentration beyond a characteristic value n_{CB} that we try to determine in the following. The distinction of IB and CB exists only when we can observe a clear gap between them. From the diagonalisation, however, all we can do is count states, as described in Sec. 4.4, and define ε_{IB} (ε_{CB}) as the energy of the highest occupied (lowest unoccupied) many-independent-electron state. By definition, then, $\varepsilon_{CB} > \varepsilon_{IB}$ and the bands would never merge, contrarily to what appears in Fig. 6.1.

We can therefore compare two alternative quantities, the energy gap $\varepsilon_{gap} = \varepsilon_{CB} - \varepsilon_{IB}$ and the average energy level spacing in the IB, Δ_{IB} . When $\varepsilon_{gap} \gg \Delta_{IB}$ the system has a clearly defined gap and cannot conduct. On the other hand,




Figure 6.2: Gap closure indicator δ_{IB} as a function of n . Error bars indicate the standard error of the mean. The area shaded in gray highlights the concentrations at which IB and CB mix. Reproduced from Carnio et al. (2017).

when $\varepsilon_{\text{gap}} \lesssim \Delta_{IB}$ we cannot distinguish a band gap among the energy differences between states, and hence there is no point in distinguishing two bands in the first place. To *qualitatively* capture the behaviour of the two bands, we then define the relative ratio between these quantities

$$\delta_{IB} = \frac{\varepsilon_{\text{gap}} - \Delta_{IB}}{\varepsilon_{\text{gap}}}, \quad (6.2)$$

such that, for $\varepsilon_{\text{gap}} \gg \Delta_{IB}$, $\delta_{IB} \sim 1$. This is indeed the case for small impurity concentrations, shown in Fig. 6.2. Upon increasing n , we observe δ_{IB} decrease and vanish for $7 \times 10^{20} \text{ cm}^{-3} \lesssim n_{CB} \lesssim 9 \times 10^{20} \text{ cm}^{-3}$, indicated by the grey area. We estimate the width of this region by eye, as δ_{IB} is only a qualitative indicator. The negative values for higher concentrations indicate that the “band gap”, intended as the separation between the highest occupied and lowest unoccupied energy levels, is smaller than the mean level spacing throughout the IB. This could be interpreted as a higher level density near ε_F , where the top of the IB merges with the CB.

In the only experiment on the MIT in Si:S we are aware of, presented in Winkler et al. (2011), the critical concentration at the Fermi energy is found to lie between 1.8 and $4.3 \times 10^{20} \text{ cm}^{-3}$. Considering the approximations at the foundation of our model, we can be satisfied with obtaining a critical concentration of the same order of magnitude. In fact, starting from a DFT treatment, we already




Figure 6.3: Shift of the average Fermi energy ε_F as a function of concentration n for the four system sizes we simulated.

know that our critical concentration is bound to be at least two times bigger: this is because DFT underestimates the correlations between electrons (see Sec. 2.2.2), and hence the average distance between impurities, which is approximately $n_c^{-1/3}$.

In Fig. 6.3 we show how the average Fermi energy shifts higher in energy with increasing concentration. After a steep increase at low concentrations, ε_F seems to stabilise and even decrease slightly. This is consistent with the picture of merging bands: when counting the states we will reach the region where the top of the IB meets the lower CB states. This behaviour, however, seems to happen for concentrations around $2n_{CB}$.

The experiments by Alexander and Holcomb (1968) on Si:P and other materials find that $n_c < n_{CB}$, namely the IB first delocalises and then merges with the CB. In our model, instead, we find the opposite. This might be due to the different material, since sulphur is a deep centre in silicon (see Sec. 2.3). Its IB forms closer to the middle of the band gap and hence might require a higher concentration to extend all the way to the CB. We need results on a semiconductor doped with shallow impurities to test this hypothesis.

6.4 Delocalisation in the impurity band

In Fig. 6.4 we add α_0 as another dimension to the DOS of Fig. 6.1. In this 2D histogram, then, we can study how states are distributed not only in energy, but also with respect to their localised character, since the exponent α_0 allows an




Figure 6.4: Distribution of the moments α_0 as a function of ϵ shifted with ϵ_F (vertical dashed line). For $N_S = 40$ we show the density plot of the distribution (from blue for low to red for high density) and the contour lines enclosing 68% (white) and 95% (black) of the α_0 's. For $N_S = 100$ we indicate the same contours (red, dashed). The Roman numbering and alternate shading indicate the regions described in the main text of Sec. 6.4. Reproduced from Carnio et al. (2017).

ordering in the localisation properties of a wave function. Using Eq. (4.13) we can show that $\alpha_0 = 3$ for the perfectly extended state $|\psi|^2 = L^{-d}$ (a plane wave), while $\alpha_0 \rightarrow \infty$ in the limit of a state localised on lattice site i , i.e. $|\psi|^2 = \delta_{i,j}$. In a plot like in Fig. 6.4, therefore, a downward shift of the points indicates delocalisation of the states.

From the data for $N_S = 40$ ($n = 4.9 \times 10^{20} \text{ cm}^{-3}$) we can distinguish these four regimes (highlighted in Fig. 6.4):

- (i) Beyond the Fermi energy $\epsilon - \epsilon_F \geq 0$ we can see an accumulation of conduction band (CB) states, with $\alpha_0 \approx 3.15$.
- (ii) For $-0.05 \text{ eV} \lesssim \epsilon - \epsilon_F \lesssim 0$, the density of states is low and α_0 increases to values around 3.4.
- (iii) The majority of the states are concentrated for energies between -0.2 eV and -0.05 eV and their α_0 remains around 3.4, with a slight decrease.

- (iv) For energies lower than -0.2 eV we see a long tail of states that are more localised the deeper we move into the IB.

Upon increasing the number of dopants to $N_S = 100$, we can see from the contour lines in Fig. 6.4 that the distinction between (i), (ii) and (iii) disappears. This representation suggests that the IB and CB not only merge with respect to energy, but with respect to the localisation properties, too. In fact, this picture supports the idea that the two bands connect first, and only later the states become delocalised. It is tempting to imagine this process as the reciprocal approach of the denser regions (i) and (iii), with the intermediate hybridisation region (ii) to form some sort of connection between the two. The information in Fig. 6.4 can shed some light on results we have seen in the previous chapters. The first observation is that some states are more delocalised than others: clearly in (iv) states are highly localised. In the plateau region (iii), instead, we expect to see some delocalised states in the shaded area where $n > n_c$. The values of α_0 for these states, however, is still significantly higher than those near ε_F . Also for $N_S = 100$, the unified band shows an increment of α_0 for lower energies. This is consistent with what we have observed in Fig. 4.4, where, for a given n_c and L , the critical state at the Fermi energy is more delocalised than that deeper in the band.

6.5 Band merging and the critical exponent

Another open question is the behaviour of ν observed in the phase diagrams of Sec. 5.6, in particular the increasing behaviour near the Fermi energy. The division of the energy axis in different regimes in Fig. 5.7 and 6.4 bears some similarities. In particular, the energies where ν increases seems to correspond to the hybridisation region (ii). Intuitively we could say that, because of the interference of the conduction band, the states near the Fermi energy are, on average, more extended. This means that, when we attach our sample to a battery, we are going to measure a higher conductivity. According to Eq. (5.2), the conductivity scales like $\sigma \sim w^\nu$ in the metallic regime. Since in the critical regime $|w| \ll 1$, an increase in conductivity can be explained by either a decrease in ν or a change in the proportionality factor.

The same argument can be formulated also for the multifractal exponent. Let us work with $l = 1$ and a certain L , like in Fig. 6.4. At this fixed λ we can use Eq.

(5.14) to say that, at a certain rescaled disorder w , α_0 scales like

$$\alpha_0(w, L) \sim a_0 + a_1 w L^{1/\nu}. \quad (6.3)$$

Here $a_0 = \alpha_0^\lambda$ is the crossing point defined in (5.12) and showed, e.g., in Fig. 5.4. In the delocalised regime, where $w > 0$, the quantity $\alpha_0(w, L) - \alpha_0^\lambda$ will be negative (cf. again Fig. 5.4), i.e. has opposite sign to w and thus a_1 is a negative constant. Equation (6.3) thus becomes

$$\alpha_0^\lambda - \alpha_0(w, L) \sim w L^{1/\nu}. \quad (6.4)$$

If, because of hybridisation with conduction band states (with $\alpha_0 \approx 3$), the average $\alpha_0(w, L)$ effectively has a lower value than expected, then for $w > 0$ we will get further away from the crossing point, and hence we will see an increase in $\alpha_0^\lambda - \alpha_0(w, L)$. If we call ν_{mix} and ν the critical exponent measured, respectively, with and without the influence of the conduction band states, then we must have

$$w L^{1/\nu_{\text{mix}}} \geq w L^{1/\nu}, \quad (6.5)$$

which immediately leads to

$$\nu_{\text{mix}} \leq \nu. \quad (6.6)$$

We therefore conclude that the hybridisation with more extended states, like near ε_F , is compatible with an effective reduction in the observed ν . Analogously, hybridising with more localised states leads to an effective increase in ν . This might explain the behaviour in the tail of localised states in (iv), where the physics of the Anderson transition seems to reemerge with ν reaching the range of its universal value, whether without (Slevin and Ohtsuki, 1999), or with interactions (Harashima and Slevin, 2014).

Chapter 7

Conclusions and outlook

In this Thesis we present an effective tight-binding model, based on an *ab initio* approach, to study the metal-insulator transition in a doped semiconductor. We start from the planning and set up of density functional theory (DFT) simulations, and work our way to the characterisation of the transition in terms of critical concentration of impurities and critical exponent. This model inherits the salient features of DFT: it describes a single electron at $T = 0$, interacting with other electrons and nuclei via the Coulomb interaction (as approximated in the exchange-correlation functional).

Since we have worked with doped silicon (Si), Ch. 2 is devoted to the DFT simulation of Si, bulk or with impurities, with phosphorous (P, Fig. 2.4) and sulphur (S, Fig. 2.5) as examples. We discuss the different treatment of single and paired impurities and the generation of disordered samples. The results of the DFT simulations – i.e. the Kohn-Sham Hamiltonians – are used to construct effective tight-binding models (ETBM) following the prescriptions given in Ch. 3. We have a system where the matrix elements of the host atom provide the background potential, and those of the impurities are an identical local perturbation, randomly placed in the lattice. This perturbation extends up to ten neighbours away from a site, and we contemplate two fundamental types: single and paired impurities (Fig. 3.2). We have shown that the spectrum of impurity states in the ETBM agrees well qualitatively with a DFT simulation of the same system (Fig. 3.4). We generate and diagonalise up to a thousand disorder realisations of the ETBM for four system sizes with different impurity concentrations (table 4.1).

Chapter 4 deals with the multifractal properties of the wave function at criticality. After a review of the main definitions, we compute the singularity spectrum

of a wave function computed with DFT and show that it has a multifractal nature (Fig. 4.1). This result is in line with previous experimental (Richardella et al., 2010), theoretical (Burmistrov et al., 2013), and numerical (Amini et al., 2014; Harashima and Slevin, 2014) studies, where the critical fluctuations of the wave function at criticality are expected to survive in the presence of the Coulomb interaction. Analogously, we compute the ensemble-averaged singularity spectrum for the ETBM near the critical concentration and show that multifractality persists also under the assumptions behind the effective models. Similarly to the non-interacting Anderson model (Rodriguez et al., 2008), the multifractal spectrum $f(\alpha)$ for the ETBM correctly respects the symmetry expected from field-theoretical models (Mirlin et al., 2006) and shows the presence of negative fractal dimensions related to rare events (Fig. 4.2). A noticeable difference, however, is that $f(\alpha)$ has a nearly-parabolic shape (weak multifractality), a common trait with the Anderson transition in $2 + \epsilon$, with $\epsilon \ll 1$, and with the power-law random banded matrix model with $b \gg 1$ (Evers and Mirlin, 2008). Finally, from the $f(\alpha)$ spectrum near the two mobility edges, we observe that the eigenstate is more delocalised near the Fermi energy than deeper in the band (Fig. 4.4). We explain this observation in Ch. 6 as the presence of different localisation regimes in the impurity band.

In Ch. 5 we study the scaling with system size and disorder of the multifractal exponents calculated in Ch. 4. We introduce the scaling theory of localisation and finite-size scaling. We check that the ETBM captures the Anderson transition by checking both the self-similarity of the distribution of the logarithmically-renormalised wave function α (see Eq. (4.8) and Fig. 5.2), and its flowing towards opposite regimes under box-size renormalisation (Fig. 5.3). After verifying the adequacy of the fitting function (Fig. 5.4), we present the phase diagrams of the critical concentration n_c and exponent ν as a function of energy from the Fermi level, $\epsilon - \epsilon_F$ (Fig. 5.7). We discuss the stability of the results for different coarse-grainings and different moments. Despite some variations in the specific values, our data consistently show some interesting trends. Over the same energy range, the critical concentration decreases to a minimum, while the critical exponent increase from 0.5 to 1. Deeper in the band, the critical concentration increases again, while the critical concentration fluctuates between values 1 and 1.5. From the concave shape of n_c we deduce that, like an Anderson band, the delocalisation transition starts in the centre of the impurity band and moves towards the tails with increasing concentration.

The behaviour of the impurity band (IB) is studied in Ch. 6. The prerequisites for a metal-insulator transition in the IB are a delocalisation transition of the Anderson type, studied in Ch. 5, and the merging with the conduction band (CB). The latter we show by computing the density of states (DOS) of the band (Fig. 6.1), and by comparing the mean level spacing to the band gap (Fig. 6.2). We then compute a 2D DOS by adding the information on the average α for each wave function (Fig. 6.4). From this plot we recognise different regimes in the band: a central accumulation of states, which becomes progressively more delocalised, a tail of highly-localised states deeper in the band, and a “hybridisation region” that connects the central peak to the conduction band. While all states undergo a delocalisation transition, these different regions might explain the variation in the critical exponent reported in the phase diagram. In particular, we show that hybridisation with extended states, e.g. from the conduction band, can lead to the observed reduction in the estimate of the critical exponent when approaching the Fermi energy.

A comparison with M. Itoh et al. (2004) is in order. As our model describes an uncompensated semiconductor, and our finite-size scaling analysis includes a large critical regime, we would expect $\nu \approx 0.5$. While this indeed happens near the Fermi energy, it does not justify the increasing trend deeper in the IB, where we find $\nu \approx 1$ to 1.5 by keeping the same critical regime. M. Itoh et al. (2004) observe nonetheless that deliberate compensation seems to lead to estimates $\nu \approx 1$. This can be reconciled with our results, since compensation can shift the Fermi energy deeper in the band, where we indeed find $\nu \gg 0.5$. As to why, in a compensated sample, $\nu \approx 1$ robustly over large critical regions, we do not have a definite answer.

In Harashima and Slevin (2014) the estimate for the critical exponent is $\nu \approx 1.3$ for the Fermi level,¹ while we find $\nu \approx 0.5$. The question they address, however, is how the Coulomb interaction changes the critical exponent of the Anderson transition. To this end, they consider a model with hydrogenic impurities (see Sec. 2.3) in an effective medium with the dielectric constant and effective mass of Si. Our model, instead, is arguably closer to the description of a doped semiconductor: we treat the host Si atoms on the same footing as the impurities, i.e. the electrons are allowed to interact also with the host nuclei and valence electrons. We also

¹They consider the ensemble of the highest-occupied Kohn-Sham eigenstates, whereas we define ϵ_F as the midpoint of the band gap. We do not expect this difference to play any significant role.

use the diamond lattice, which has a lower connectivity ($Z = 4$), compared to the three-dimensional simple cubic lattice ($Z = 6$). This should change the value of the critical disorder, as already observed by Anderson (1958), but it is not known to change the critical exponent (see, e.g. Puschmann et al. (2015) for a study of Anderson localisation on lattices with random connectivities). Another more technical difference is that the effective models are *not* the self-consistent solution of the Kohn-Sham Hamiltonians. They are rather an approximation, as discussed in Sec. 3.2. Still, since this approximation works quite well in reproducing the DFT results (see Fig. 3.4), we do not expect it to yield such a large discrepancy in the estimates of ν . Finally, at least in the case of sulphur, we only include the local strain fields encoded in the catalogues, rather than allowing the geometry relaxation of the whole system when a significant number of dopants is present. We argue, however, that this systematic error is negligible: the deformation of the lattice extends, at most, to the first few shells of nearest neighbours, as seen in Fig. 2.7. Its correction, moreover, would preclude the possibility to build catalogues, and hence effective model, so the disorder average would come at the price of several DFT simulations. All considered, then, our results do not contradict those of Harashima and Slevin (2014), and, in fact, $\nu \approx 1.3$ is compatible with the values we find deeper in the IB. While they also compute the DOS and show two bands merging, it would be interesting to have a complete picture with a phase diagram like Fig. 5.7 and a 2D histogram like Fig. 6.4 to study how the two bands interact.

Another factor that might change the value of the critical exponent is the correlation in the disorder potential. Contrarily to the Anderson model, where the on-diagonal elements are random variables, our Hamiltonians consists of a finite set of matrix blocks (from the catalogues) appearing in random rows and columns, making our disorder correlated. From the extended Harris criterion (Weinrib and Halperin, 1983), however, we expect the critical exponent to be *higher* or equal the uncorrelated case $\nu \approx 1.59$ (Ndawana et al., 2004). For this reason, a plausible picture is that the universality class of the Anderson transition with the Coulomb interaction is characterised by $\nu \approx 1.3$ (Harashima and Slevin, 2014), where, in a real material, the influence of the conduction band causes ν to decrease to roughly 0.5 near the Fermi energy. To test this hypothesis we need to isolate this band hybridisation process. Ideally, we could start from the band of an Anderson model, where a delocalisation-localisation transition happens like in the impurity band, and mix it with a band of delocalised states, mimicking the conduction band of a semiconductor. A tight-binding Hamiltonian that describes the merging of an

impurity and a conduction band can be found, e.g., in Radjenovic and Tjapkin (1989). In Plyushchay et al. (2003), instead, the transition in a two-band Anderson model is studied, but not in the case that a delocalising “impurity” band mixes with an extended “conduction” band. If our deduction in Sec. 6.5 is correct, we should see an exponent $\nu_0 \approx 1.59$ (Rodriguez et al., 2011) when the impurity band delocalises without the influence of the conduction band, and $\nu \leq \nu_0$ otherwise. If this is the case, we could exclude the Coulomb interaction as a *direct* cause of the reduction of ν , although still involved in the formation and merging of the bands.

The estimate of ν from the two-band Anderson model could also be useful in establishing whether the estimate $\nu \approx 0.5$ at the Fermi energy is a band-related effect. As mentioned in Ch. 1 for the experimental results, this specific value is interesting because it violates the Harris criterion $\nu > 2/d$, where $d = 3$ is the dimensionality of the space. This bound on ν was proved as a general theorem by Chayes et al. (1986), assuming the divergence at the critical point of an “appropriately defined” correlation length. This correlation length essentially corresponds to the localisation length of Eq. (3.6), namely the localised phase must be characterised by an exponentially suppressed probability of finding an electron at long distances (refer to Chayes et al. (1986) for a formal definition). In the hybridisation regime, then, this requirement may not be met any more, since the ensemble of states we consider includes both impurity and conduction band states with different localisation properties (see Sec. 6.4 and Fig. 6.4). Clearly further research is needed to understand this hybridisation process and its impact on the critical exponent.

The effective reduction of ν near the Fermi energy might not be (solely) a band effect. Ideally we would like to analyse how the picture depends on the Coulomb interaction. Tuning the strength of the interaction seems reminiscent of the approach of Harashima and Slevin (2014), who renormalise the potentials in Eq. 2.7 with the effective parameters of silicon. We notice, however, that the Coulomb interaction regulates also the attractions between nuclei and core electrons, the formation of a lattice and, ultimately, the stability of matter! Simulating an alternative reality with a different electromagnetic interaction sounds like a fascinating endeavour, one that should be handled with care, especially considered the lack of experimental results to compare potential outcomes to.

The critical exponent of the transition might also depend on material-specific effects, in our case on properties of Si:S, while most studies report on semiconductors doped with shallow impurities, in particular Si:P. When we introduce

many P impurities, the shallow energy level of one atom broadens into a band, which, compared to the case of a deep-level donor like sulphur, will be closer to the conduction band. In Sec. 2.3 we found that a P atom in Si introduces an impurity level around -0.016 eV, while a S atom around -0.150 eV (compare it with the DOS of Fig. 6.1). What happens, then, to the hybridisation region of Fig. 6.4 in Si:P? Since the P level is 10 times closer to the CB than the S level (see Fig. 2.6), we imagine the behaviour of Fig. 6.4 to occur on an energy scale 10 times smaller, which implies that we would still see a reduction of the critical exponent due to band hybridisation.

Our method can be adapted also to other systems, in order to potentially discriminate material-specific effects. The essential step consists in writing a tight-binding model that suits the material and identifying the essential blocks that constitute it. For binary materials, like Si:B, Si:As, Si:Sb, Ge:Sb, or silicon doped with other chalcogens, the procedure is similar to that presented in Ch. 3. If more species are involved, for instance in the case of co-doping (Si:P,B or Ge:Ga,As), the number of configurations to include in a catalogue will grow, but the fundamental principles in the construction will not change. We would then run DFT simulations to extract the matrix terms that describe such blocks and use them to build an effective model. If the effective models correctly reproduce the properties of the DFT simulations, we can proceed with studying a MIT in said material. Of course, the feasibility of such a project also depends on the size of the smallest system we need to compute, as discussed for Si:P in Sec. 3.5. The treatment with effective tight-binding models is not limited to three-dimensional materials. Density functional theory is routinely used to characterise two-dimensional nanostructures created in the laboratory, e.g. Wilson et al. (2017). With an effective tight-binding model we can then sample more efficiently the ensemble of different configurations. This also opens the possibility of an independent confirmation of the interaction-enabled MIT in 2D, beyond the theoretical study of Punnoose and Finkel'stein (2005).

Further investigation should also include studying larger system sizes, for two reasons. Firstly, $\nu \approx 0.5$ might result from taking a broad critical region, as described in the work of M. Itoh et al. (2004). This observation, however, is not consistent with the results deeper in the IB, where we take similarly broad critical regions and obtain $\nu \approx 1$ to 1.5. New numerical evidence on larger systems would clarify the matter. As a second reason, larger systems lead to more accurate estimation of the critical parameters, as discussed in Ch. 5. A

possibility to circumvent the technical difficulties related to the diagonalisation of large systems, then, is the “Kernel Polynomial Method” (KPM), an alternative way to characterise the transition. This method has been successfully employed in Jung et al. (2012, 2016) to study the Anderson MIT with magnetic impurities. In the case of vanishing exchange interaction, the non-interacting Anderson case, a finite-size scaling analysis yields $\nu = 1.48 \pm 0.06$. In Jung et al. (2012) the authors claim that the core of the KPM consists of matrix multiplications, which use little memory and scale as $O(N)$ when the Hamiltonian matrix is in a sparse representation. Implementation and optimisation on parallel architectures is also under research (Zhang et al., 2013; Kreutzer et al., 2015). Nevertheless, the presence of the overlap matrix must be taken into account and implemented in the method, as explained e.g. in Röder et al. (1997), and might make this method less convenient.

In conclusion, we have constructed an *ab initio* model of a doped semiconductor (Si:S) that replicates the essential processes of the metal-insulator transition in the impurity band, namely the hybridisation with the conduction band and the delocalisation of its states. To the best of our knowledge, this is the first time Anderson localisation is observed in an *ab initio* model of a doped semiconductor. Because of its complexity, the investigation is still in its early stages in terms of sophistication and accuracy, similarly to the studies on the Anderson model in the 1980s compared to now. The solution of the generalised eigenvalue problem of Eq. (3.9) is certainly the hardest numerical challenge and strongly limits the reachable system sizes. Despite this, the critical exponents we obtain seem to bridge the gap between experimental observations and theoretical estimates. This raises a number of questions, delineated above, aimed at better understanding and characterising the delocalisation transition, especially in relation to the hybridisation of the impurity and conduction bands. The connection to the world of material science, incorporated in our method, represents a paradigm shift that holds the promise of new insights on, and potentially resolution of, questions on Anderson localisation in real materials.

Appendix A

Fit results

In this Appendix we present tables of the stable and robust fits that constitute the basis of the pictures in Ch. 5. The values of q and system sizes are indicated in the captions, with results listed by decreasing energy and coarse-graining λ . The concentrations used for each energy lie in the interval (n_0, n_f) . Energies are expressed in eV and all concentrations in units of 10^{20} cm^{-3} . Uncertainties on the critical exponent ν and concentration n_c are 95% confidence intervals. N_p and N_D indicate, respectively, the number of parameters and of data points used, while χ^2 and p are the values of the χ^2 statistics and the goodness-of-fit probability. The expansion is in the order m_L, m_ρ .

Table A.1: Fit results for $L \leq 20$ and $q = 0$.

E_i	λ	$(n_0; n_f)$	ν	n_c	N_p	N_D	χ^2	p	$m_L m_\rho$
0.000	1/2	(6.2; 14.6)	0.60 ± 0.18	10.39 ± 0.76	6	60(0.11)	47	0.73	3 1
-0.011	1/2	(6.2; 14.6)	0.63 ± 0.17	10.57 ± 0.78	6	60(0.12)	45	0.80	3 1
-0.023	1/2	(5.6; 16.8)	0.68 ± 0.11	11.34 ± 0.71	7	65(0.13)	69	0.14	3 2
-0.034	1/2	(6.8; 16.0)	0.90 ± 0.27	11.58 ± 0.99	6	57(0.13)	63	0.11	3 1
-0.045	1/2	(6.2; 14.6)	0.88 ± 0.22	10.44 ± 0.75	6	60(0.13)	48	0.70	3 1
-0.057	1/2	(5.7; 13.3)	0.83 ± 0.21	9.60 ± 0.66	6	61(0.12)	55	0.48	3 1
-0.068	1/2	(5.0; 11.8)	0.88 ± 0.25	8.34 ± 0.60	6	64(0.13)	54	0.62	3 1
-0.079	1/2	(4.0; 9.4)	0.90 ± 0.25	6.65 ± 0.41	6	74(0.13)	57	0.83	3 1
-0.102	1/2	(0.4; 1.2)	1.04 ± 0.41	0.78 ± 0.09	6	31(0.47)	16	0.92	3 1
-0.113	1/2	(0.5; 1.3)	0.74 ± 0.24	0.85 ± 0.07	6	29(0.50)	23	0.49	3 1
-0.125	1/2	(0.6; 1.4)	0.67 ± 0.24	0.94 ± 0.07	6	30(0.53)	20	0.71	3 1
-0.136	1/2	(0.7; 1.7)	0.82 ± 0.29	1.15 ± 0.09	6	37(0.51)	30	0.52	3 1
-0.147	1/2	(0.8; 2.0)	2.06 ± 1.31	1.41 ± 0.20	6	38(0.55)	32	0.45	3 1
-0.159	1/2	(1.1; 2.5)	1.34 ± 0.50	1.78 ± 0.15	6	40(0.45)	37	0.31	3 1
-0.170	1/2	(1.3; 3.1)	1.20 ± 0.37	2.15 ± 0.15	6	40(0.41)	40	0.21	3 1
-0.181	1/2	(1.8; 3.4)	1.11 ± 0.44	2.54 ± 0.17	6	31(0.33)	18	0.85	3 1

Table A.1 continued from previous page

E_i	λ	$(n_0; n_f)$	ν	n_c	N_p	N_D (prec.)	χ^2	p	$m_L m_\rho$
-0.193	1/2	(2.3; 4.3)	2.14 ± 1.17	3.27 ± 0.34	6	38(0.28)	40	0.15	3 1
-0.204	1/2	(2.6; 4.8)	1.57 ± 0.63	3.54 ± 0.26	6	44(0.28)	52	0.07	3 1
-0.215	1/2	(3.1; 5.7)	1.47 ± 0.51	4.42 ± 0.28	6	54(0.26)	48	0.46	3 1
-0.227	1/2	(3.4; 6.2)	1.65 ± 0.53	4.93 ± 0.30	6	61(0.25)	59	0.33	3 1
-0.238	1/2	(4.2; 7.8)	1.33 ± 0.37	5.77 ± 0.26	6	72(0.21)	61	0.64	3 1
-0.249	1/2	(4.7; 8.7)	1.50 ± 0.46	6.77 ± 0.39	6	72(0.21)	69	0.38	3 1
-0.261	1/2	(5.3; 9.8)	1.66 ± 0.69	7.46 ± 0.57	6	65(0.21)	69	0.18	3 1
-0.272	1/2	(6.1; 11.3)	1.57 ± 0.70	8.60 ± 0.77	6	58(0.20)	36	0.96	3 1
-0.283	1/2	(6.3; 11.7)	2.49 ± 1.59	9.02 ± 1.27	6	55(0.21)	49	0.47	3 1
-0.295	1/2	(6.0; 11.2)	1.49 ± 0.58	8.53 ± 0.67	6	58(0.25)	47	0.66	3 1
-0.306	1/2	(6.4; 12.0)	1.73 ± 0.78	9.38 ± 0.96	6	56(0.25)	65	0.07	3 1
-0.317	1/2	(6.7; 12.4)	1.56 ± 0.69	9.39 ± 0.87	6	52(0.26)	46	0.46	3 1
-0.329	1/2	(6.9; 12.9)	2.09 ± 1.07	10.03 ± 1.16	6	53(0.27)	38	0.82	3 1
0.000	1/4	(6.8; 12.6)	0.48 ± 0.16	9.71 ± 0.60	6	52(0.10)	46	0.49	3 1
-0.011	1/4	(7.3; 13.5)	0.47 ± 0.11	10.33 ± 0.43	6	51(0.10)	42	0.59	3 1
-0.023	1/4	(7.5; 13.9)	0.57 ± 0.12	10.66 ± 0.39	6	49(0.09)	54	0.12	3 1
-0.034	1/4	(7.0; 13.0)	0.69 ± 0.12	10.08 ± 0.37	6	51(0.09)	52	0.22	3 1
-0.045	1/4	(6.7; 12.4)	0.76 ± 0.14	9.43 ± 0.35	6	52(0.09)	51	0.30	3 1
-0.057	1/4	(5.9; 11.1)	0.79 ± 0.14	8.36 ± 0.27	6	58(0.09)	38	0.92	3 1
-0.068	1/4	(4.9; 9.1)	0.87 ± 0.15	6.94 ± 0.20	6	70(0.09)	77	0.12	3 1
-0.079	1/4	(3.8; 7.0)	0.95 ± 0.15	5.33 ± 0.15	6	69(0.08)	62	0.50	3 1
-0.091	1/4	(2.6; 4.8)	1.33 ± 0.43	3.68 ± 0.22	6	44(0.09)	50	0.10	3 1
-0.136	1/4	(1.8; 3.4)	1.05 ± 0.32	2.56 ± 0.14	6	31(0.10)	19	0.80	3 1
-0.147	1/4	(1.9; 3.5)	1.74 ± 0.64	2.88 ± 0.24	6	33(0.12)	37	0.10	3 1
-0.159	1/4	(2.1; 3.9)	1.53 ± 0.54	3.01 ± 0.20	6	35(0.12)	35	0.19	3 1
-0.170	1/4	(2.5; 4.7)	1.20 ± 0.32	3.59 ± 0.18	6	41(0.11)	43	0.16	3 1
-0.181	1/4	(2.8; 5.2)	1.21 ± 0.28	4.06 ± 0.16	7	47(0.12)	54	0.07	3 2
-0.193	1/4	(3.0; 6.2)	0.95 ± 0.14	4.49 ± 0.14	6	66(0.11)	82	0.03	3 1
-0.204	1/4	(3.6; 6.6)	1.03 ± 0.17	5.12 ± 0.16	6	67(0.11)	66	0.30	3 1
-0.215	1/4	(4.1; 7.7)	0.96 ± 0.14	5.66 ± 0.15	6	73(0.11)	87	0.05	3 1
-0.227	1/4	(4.4; 8.2)	1.09 ± 0.19	6.29 ± 0.19	6	72(0.12)	85	0.06	3 1
-0.238	1/4	(4.9; 9.1)	1.03 ± 0.19	6.99 ± 0.23	6	70(0.12)	68	0.34	3 1
-0.249	1/4	(5.3; 9.9)	1.08 ± 0.19	7.80 ± 0.29	7	66(0.13)	70	0.15	3 2
-0.261	1/4	(5.9; 10.9)	1.19 ± 0.28	8.38 ± 0.39	6	59(0.13)	66	0.11	3 1
-0.272	1/4	(6.3; 11.7)	1.05 ± 0.21	8.94 ± 0.36	6	55(0.13)	54	0.29	3 1
-0.283	1/4	(6.8; 12.6)	1.05 ± 0.20	9.65 ± 0.39	6	52(0.14)	51	0.30	3 1
-0.295	1/4	(6.7; 12.5)	1.17 ± 0.22	9.56 ± 0.38	6	55(0.14)	41	0.78	3 1
-0.306	1/4	(7.3; 13.5)	0.94 ± 0.17	10.16 ± 0.36	6	51(0.15)	53	0.20	3 1
-0.317	1/4	(7.5; 13.9)	1.16 ± 0.25	10.66 ± 0.48	6	49(0.15)	49	0.26	3 1
-0.329	1/4	(7.7; 14.3)	1.22 ± 0.29	10.96 ± 0.52	6	47(0.16)	35	0.73	3 1
-0.340	1/4	(8.0; 15.0)	1.36 ± 0.40	11.47 ± 0.65	6	46(0.17)	38	0.57	3 1
-0.351	1/4	(8.3; 15.3)	1.15 ± 0.28	11.84 ± 0.58	6	47(0.18)	35	0.73	3 1

Table A.1 continued from previous page

E_i	λ	$(n_0; n_f)$	ν	n_c	N_p	N_D (prec.)	χ^2	p	$m_L m_\rho$
-0.363	1/4	(8.8; 16.3)	1.18 ± 0.33	12.46 ± 0.67	6	44(0.18)	31	0.80	3 1
-0.374	1/4	(9.4; 17.4)	1.44 ± 0.49	13.43 ± 0.96	6	43(0.20)	31	0.76	3 1

Table A.2: Fit results for $L \leq 20$ and $q = 1$.

E_i	λ	$(n_0; n_f)$	ν	n_c	N_p	N_D	χ^2	p	$m_L m_\rho$
0.000	1/2	(7.3; 13.5)	0.59 ± 0.25	10.25 ± 0.83	6	51(0.12)	38	0.77	3 1
-0.011	1/2	(7.4; 13.7)	0.48 ± 0.17	10.51 ± 0.68	6	51(0.13)	37	0.80	3 1
-0.023	1/2	(7.8; 14.6)	0.62 ± 0.27	11.21 ± 0.84	6	48(0.14)	52	0.14	3 1
-0.034	1/2	(7.6; 14.0)	0.83 ± 0.35	10.81 ± 0.94	6	49(0.14)	51	0.19	3 1
-0.045	1/2	(7.1; 13.3)	0.75 ± 0.26	10.20 ± 0.72	6	52(0.14)	58	0.11	3 1
-0.057	1/2	(6.8; 12.6)	1.01 ± 0.44	9.66 ± 0.86	6	52(0.14)	57	0.12	3 1
-0.068	1/2	(5.7; 10.5)	0.76 ± 0.28	8.11 ± 0.56	6	61(0.15)	56	0.44	3 1
-0.079	1/2	(4.9; 9.1)	0.74 ± 0.26	6.91 ± 0.37	6	70(0.14)	49	0.91	3 1
-0.102	1/2	(0.4; 1.1)	1.03 ± 0.51	0.65 ± 0.11	6	26(0.73)	10	0.98	3 1
-0.113	1/2	(0.4; 1.1)	0.82 ± 0.41	0.75 ± 0.09	6	26(1.09)	23	0.28	3 1
-0.125	1/2	(0.5; 1.4)	0.69 ± 0.25	0.86 ± 0.09	6	35(0.98)	18	0.94	3 1
-0.136	1/2	(0.7; 1.5)	1.07 ± 0.56	0.99 ± 0.16	6	35(0.87)	26	0.61	3 1
-0.147	1/2	(0.7; 2.0)	1.63 ± 0.79	1.25 ± 0.17	6	45(1.13)	46	0.21	3 1
-0.159	1/2	(1.0; 2.2)	1.37 ± 0.66	1.59 ± 0.16	6	39(0.84)	41	0.16	3 1
-0.170	1/2	(1.2; 2.8)	1.33 ± 0.58	2.01 ± 0.19	6	40(0.71)	42	0.16	3 1
-0.181	1/2	(1.4; 3.4)	1.23 ± 0.43	2.37 ± 0.17	6	41(0.60)	27	0.82	3 1
-0.193	1/2	(1.5; 4.5)	1.71 ± 0.51	3.01 ± 0.24	6	59(0.58)	69	0.07	3 1
-0.204	1/2	(2.0; 4.8)	1.50 ± 0.46	3.41 ± 0.24	6	52(0.46)	61	0.06	3 1
-0.215	1/2	(3.0; 5.6)	1.74 ± 0.74	4.29 ± 0.33	6	52(0.33)	50	0.31	3 1
-0.227	1/2	(3.3; 6.1)	1.58 ± 0.57	4.65 ± 0.30	6	60(0.33)	53	0.49	3 1
-0.238	1/2	(4.0; 7.4)	1.76 ± 0.60	5.78 ± 0.36	6	71(0.27)	69	0.36	3 1
-0.249	1/2	(4.6; 8.6)	1.37 ± 0.44	6.50 ± 0.35	6	71(0.26)	78	0.14	3 1
-0.261	1/2	(5.3; 9.9)	1.79 ± 0.88	7.76 ± 0.75	6	66(0.25)	72	0.14	3 1
-0.272	1/2	(6.1; 11.3)	1.62 ± 0.77	8.53 ± 0.81	6	58(0.24)	40	0.88	3 1
-0.283	1/2	(6.4; 11.8)	1.93 ± 1.08	9.01 ± 1.05	6	55(0.26)	55	0.27	3 1
-0.295	1/2	(6.0; 11.2)	1.58 ± 0.70	8.43 ± 0.73	6	58(0.30)	54	0.39	3 1
-0.306	1/2	(6.2; 11.4)	1.32 ± 0.53	8.74 ± 0.70	6	56(0.32)	66	0.06	3 1
-0.317	1/2	(6.6; 12.2)	1.45 ± 0.63	9.28 ± 0.82	6	54(0.33)	43	0.66	3 1
-0.329	1/2	(6.9; 12.7)	2.38 ± 1.49	9.79 ± 1.34	6	52(0.34)	34	0.90	3 1
0.000	1/4	(6.8; 12.6)	0.53 ± 0.18	9.71 ± 0.61	6	52(0.13)	45	0.50	3 1
-0.011	1/4	(7.2; 13.4)	0.67 ± 0.22	10.28 ± 0.62	6	49(0.13)	30	0.94	3 1
-0.023	1/4	(7.4; 13.8)	0.66 ± 0.16	10.52 ± 0.47	6	48(0.13)	46	0.29	3 1
-0.034	1/4	(7.0; 13.0)	0.70 ± 0.13	10.05 ± 0.39	6	51(0.13)	45	0.46	3 1
-0.045	1/4	(6.7; 12.4)	0.74 ± 0.15	9.49 ± 0.38	6	52(0.13)	49	0.35	3 1
-0.057	1/4	(5.9; 11.1)	0.80 ± 0.17	8.42 ± 0.33	6	58(0.13)	37	0.94	3 1
-0.068	1/4	(4.8; 9.0)	0.84 ± 0.16	6.83 ± 0.21	6	71(0.13)	73	0.23	3 1

Table A.2 continued from previous page

E_i	λ	$(n_0; n_f)$	ν	n_c	N_p	N_D (prec.)	χ^2	p	$m_L m_\rho$
-0.079	1/4	(2.8; 6.6)	1.48 ± 0.33	4.65 ± 0.26	6	70(0.14)	64	0.49	3 1
-0.091	1/4	(1.1; 2.5)	2.06 ± 0.97	1.70 ± 0.21	6	40(0.21)	41	0.18	3 1
-0.102	1/4	(0.7; 1.7)	1.07 ± 0.35	1.21 ± 0.10	6	37(0.27)	28	0.60	3 1
-0.113	1/4	(0.7; 1.7)	1.06 ± 0.32	1.10 ± 0.08	6	37(0.31)	45	0.05	3 1
-0.125	1/4	(0.7; 1.7)	0.97 ± 0.26	1.19 ± 0.07	6	37(0.37)	36	0.25	3 1
-0.136	1/4	(0.9; 2.1)	1.39 ± 0.49	1.43 ± 0.11	6	39(0.36)	26	0.79	3 1
-0.147	1/4	(1.0; 2.4)	2.10 ± 1.02	1.68 ± 0.19	6	40(0.40)	43	0.14	3 1
-0.159	1/4	(1.2; 2.8)	1.62 ± 0.53	1.96 ± 0.16	6	40(0.39)	46	0.08	3 1
-0.170	1/4	(1.6; 3.6)	1.51 ± 0.39	2.57 ± 0.16	6	43(0.30)	38	0.44	3 1
-0.181	1/4	(1.8; 4.2)	1.77 ± 0.51	2.97 ± 0.20	6	46(0.28)	39	0.50	3 1
-0.204	1/4	(3.0; 5.6)	1.29 ± 0.32	4.38 ± 0.21	6	52(0.22)	51	0.28	3 1
-0.215	1/4	(3.5; 6.5)	1.23 ± 0.26	5.00 ± 0.19	6	65(0.20)	55	0.62	3 1
-0.238	1/4	(4.5; 8.3)	1.17 ± 0.24	6.32 ± 0.21	6	73(0.20)	55	0.85	3 1
-0.249	1/4	(4.9; 9.1)	1.30 ± 0.31	7.06 ± 0.31	6	70(0.21)	74	0.19	3 1
-0.261	1/4	(5.6; 10.4)	1.36 ± 0.40	7.90 ± 0.43	6	62(0.20)	67	0.16	3 1
-0.272	1/4	(6.0; 11.2)	1.41 ± 0.40	8.56 ± 0.50	6	58(0.21)	62	0.16	3 1
-0.283	1/4	(6.4; 12.0)	1.40 ± 0.35	9.20 ± 0.50	6	56(0.21)	60	0.15	3 1
-0.295	1/4	(6.2; 11.6)	1.25 ± 0.29	9.05 ± 0.46	6	57(0.24)	37	0.93	3 1
-0.306	1/4	(6.9; 12.7)	1.43 ± 0.39	9.88 ± 0.59	6	52(0.24)	49	0.34	3 1
-0.317	1/4	(7.0; 13.0)	1.52 ± 0.42	10.08 ± 0.63	6	51(0.25)	51	0.25	3 1
-0.329	1/4	(7.4; 13.7)	1.71 ± 0.58	10.47 ± 0.74	6	51(0.26)	38	0.77	3 1
-0.340	1/4	(7.4; 13.7)	1.46 ± 0.42	10.53 ± 0.63	6	51(0.28)	50	0.28	3 1
-0.351	1/4	(7.9; 14.7)	1.48 ± 0.48	11.25 ± 0.72	6	47(0.29)	36	0.69	3 1
-0.363	1/4	(8.2; 15.2)	1.29 ± 0.41	11.63 ± 0.70	6	47(0.30)	38	0.62	3 1
-0.374	1/4	(8.9; 16.5)	2.00 ± 1.05	12.65 ± 1.35	6	44(0.31)	31	0.79	3 1

Table A.3: Fit results for $L \leq 22$ and $q = 0$.

E_i	λ	$(n_0; n_f)$	ν	n_c	N_p	N_D	χ^2	p	$m_L m_\rho$
0.000	1/2	(7.7; 14.3)	0.49 ± 0.13	10.99 ± 0.44	6	72(0.11)	57	0.78	3 1
-0.011	1/2	(7.7; 14.3)	0.58 ± 0.14	11.01 ± 0.44	6	72(0.11)	62	0.61	3 1
-0.023	1/2	(8.3; 13.8)	0.68 ± 0.19	11.01 ± 0.43	6	58(0.11)	65	0.11	3 1
-0.034	1/2	(7.8; 13.0)	0.73 ± 0.18	10.38 ± 0.37	6	62(0.12)	59	0.36	3 1
-0.045	1/2	(7.3; 12.1)	0.75 ± 0.17	9.69 ± 0.33	6	68(0.13)	81	0.05	3 1
-0.057	1/2	(6.8; 10.2)	0.74 ± 0.23	8.60 ± 0.34	6	68(0.13)	80	0.06	3 1
-0.068	1/2	(6.2; 9.3)	0.81 ± 0.25	7.79 ± 0.30	6	84(0.13)	94	0.11	3 1
-0.079	1/2	(5.1; 7.7)	0.85 ± 0.33	6.43 ± 0.26	6	90(0.12)	88	0.35	3 1
-0.102	1/2	(0.2; 1.4)	1.01 ± 0.26	0.82 ± 0.08	6	45(0.43)	27	0.92	3 1
-0.113	1/2	(0.2; 1.4)	0.77 ± 0.16	0.82 ± 0.05	6	45(0.55)	39	0.46	3 1
-0.125	1/2	(0.4; 1.6)	0.69 ± 0.13	1.00 ± 0.06	6	52(0.69)	53	0.21	3 1
-0.136	1/2	(0.6; 1.8)	0.91 ± 0.23	1.21 ± 0.08	6	50(0.59)	56	0.10	3 1
-0.147	1/2	(0.8; 2.3)	1.34 ± 0.39	1.48 ± 0.11	6	54(0.59)	63	0.07	3 1

Table A.3 continued from previous page

E_i	λ	$(n_0; n_f)$	ν	n_c	N_p	N_D (prec.)	χ^2	p	$m_L m_\rho$
-0.159	1/2	(1.1; 2.5)	1.26 ± 0.35	1.84 ± 0.12	6	47(0.46)	48	0.22	3 1
-0.170	1/2	(1.3; 3.1)	1.16 ± 0.27	2.21 ± 0.12	6	49(0.42)	45	0.38	3 1
-0.181	1/2	(1.6; 3.6)	1.54 ± 0.40	2.65 ± 0.16	6	53(0.38)	45	0.57	3 1
-0.193	1/2	(2.3; 4.3)	1.73 ± 0.57	3.29 ± 0.20	6	47(0.29)	53	0.10	3 1
-0.204	1/2	(2.3; 4.9)	1.54 ± 0.36	3.65 ± 0.18	6	61(0.32)	71	0.07	3 1
-0.215	1/2	(3.1; 5.7)	1.46 ± 0.37	4.43 ± 0.21	6	67(0.27)	52	0.79	3 1
-0.227	1/2	(3.3; 6.8)	1.49 ± 0.28	5.06 ± 0.20	6	102(0.24)	119	0.06	3 1
-0.238	1/2	(3.9; 8.1)	1.58 ± 0.29	5.90 ± 0.21	6	127(0.21)	131	0.25	3 1
-0.249	1/2	(5.0; 8.4)	1.25 ± 0.30	6.72 ± 0.23	6	106(0.20)	100	0.48	3 1
-0.261	1/2	(6.1; 9.1)	1.32 ± 0.43	7.68 ± 0.33	6	86(0.20)	86	0.31	3 1
-0.272	1/2	(5.4; 11.2)	1.68 ± 0.33	8.32 ± 0.38	6	131(0.22)	140	0.17	3 1
-0.283	1/2	(6.1; 11.3)	1.43 ± 0.30	8.62 ± 0.35	6	107(0.22)	97	0.60	3 1
-0.295	1/2	(6.6; 11.0)	1.33 ± 0.34	8.83 ± 0.38	6	84(0.24)	69	0.75	3 1
-0.317	1/2	(8.0; 12.0)	1.25 ± 0.46	9.97 ± 0.46	6	52(0.24)	51	0.30	3 1
0.000	1/2	(7.7; 14.3)	0.49 ± 0.13	10.99 ± 0.44	6	72(0.11)	57	0.78	3 1
-0.011	1/2	(7.7; 14.3)	0.58 ± 0.14	11.01 ± 0.44	6	72(0.11)	62	0.61	3 1
-0.023	1/2	(8.3; 13.8)	0.68 ± 0.19	11.01 ± 0.43	6	58(0.11)	65	0.11	3 1
-0.034	1/2	(7.8; 13.0)	0.73 ± 0.18	10.38 ± 0.37	6	62(0.12)	59	0.36	3 1
-0.045	1/2	(7.3; 12.1)	0.75 ± 0.17	9.69 ± 0.33	6	68(0.13)	81	0.05	3 1
-0.057	1/2	(6.8; 10.2)	0.74 ± 0.23	8.60 ± 0.34	6	68(0.13)	80	0.06	3 1
-0.068	1/2	(6.2; 9.3)	0.81 ± 0.25	7.79 ± 0.30	6	84(0.13)	94	0.11	3 1
-0.079	1/2	(5.1; 7.7)	0.85 ± 0.33	6.43 ± 0.26	6	90(0.12)	88	0.35	3 1
-0.102	1/2	(0.2; 1.4)	1.01 ± 0.26	0.82 ± 0.08	6	45(0.43)	27	0.92	3 1
-0.113	1/2	(0.2; 1.4)	0.77 ± 0.16	0.82 ± 0.05	6	45(0.55)	39	0.46	3 1
-0.125	1/2	(0.4; 1.6)	0.69 ± 0.13	1.00 ± 0.06	6	52(0.69)	53	0.21	3 1
-0.136	1/2	(0.6; 1.8)	0.91 ± 0.23	1.21 ± 0.08	6	50(0.59)	56	0.10	3 1
-0.147	1/2	(0.8; 2.3)	1.34 ± 0.39	1.48 ± 0.11	6	54(0.59)	63	0.07	3 1
-0.159	1/2	(1.1; 2.5)	1.26 ± 0.35	1.84 ± 0.12	6	47(0.46)	48	0.22	3 1
-0.170	1/2	(1.3; 3.1)	1.16 ± 0.27	2.21 ± 0.12	6	49(0.42)	45	0.38	3 1
-0.181	1/2	(1.6; 3.6)	1.54 ± 0.40	2.65 ± 0.16	6	53(0.38)	45	0.57	3 1
-0.193	1/2	(2.3; 4.3)	1.73 ± 0.57	3.29 ± 0.20	6	47(0.29)	53	0.10	3 1
-0.204	1/2	(2.3; 4.9)	1.54 ± 0.36	3.65 ± 0.18	6	61(0.32)	71	0.07	3 1
-0.215	1/2	(3.1; 5.7)	1.46 ± 0.37	4.43 ± 0.21	6	67(0.27)	52	0.79	3 1
-0.227	1/2	(3.3; 6.8)	1.49 ± 0.28	5.06 ± 0.20	6	102(0.24)	119	0.06	3 1
-0.238	1/2	(3.9; 8.1)	1.58 ± 0.29	5.90 ± 0.21	6	127(0.21)	131	0.25	3 1
-0.249	1/2	(5.0; 8.4)	1.25 ± 0.30	6.72 ± 0.23	6	106(0.20)	100	0.48	3 1
-0.261	1/2	(6.1; 9.1)	1.32 ± 0.43	7.68 ± 0.33	6	86(0.20)	86	0.31	3 1
-0.272	1/2	(5.4; 11.2)	1.68 ± 0.33	8.32 ± 0.38	6	131(0.22)	140	0.17	3 1
-0.283	1/2	(6.1; 11.3)	1.43 ± 0.30	8.62 ± 0.35	6	107(0.22)	97	0.60	3 1
-0.295	1/2	(6.6; 11.0)	1.33 ± 0.34	8.83 ± 0.38	6	84(0.24)	69	0.75	3 1
-0.317	1/2	(8.0; 12.0)	1.25 ± 0.46	9.97 ± 0.46	6	52(0.24)	51	0.30	3 1
0.000	1/4	(7.5; 13.9)	0.41 ± 0.06	10.68 ± 0.25	6	76(0.09)	71	0.43	3 1

Table A.3 continued from previous page

E_i	λ	$(n_0; n_f)$	ν	n_c	N_p	N_D (prec.)	χ^2	p	$m_L m_\rho$
-0.011	1/4	(7.4; 13.8)	0.48 ± 0.07	10.57 ± 0.22	6	74(0.09)	77	0.21	3 1
-0.023	1/4	(8.5; 12.7)	0.68 ± 0.18	10.56 ± 0.30	6	46(0.09)	53	0.09	3 1
-0.034	1/4	(6.9; 12.9)	0.65 ± 0.06	9.98 ± 0.24	7	85(0.09)	96	0.08	3 2
-0.045	1/4	(6.8; 11.3)	0.72 ± 0.09	9.09 ± 0.18	6	80(0.09)	95	0.05	3 1
-0.057	1/4	(6.0; 10.0)	0.72 ± 0.09	7.98 ± 0.14	6	99(0.09)	110	0.10	3 1
-0.068	1/4	(4.6; 8.5)	0.89 ± 0.12	6.75 ± 0.16	6	64(0.09)	74	0.07	3 1
-0.159	1/4	(1.7; 4.1)	1.20 ± 0.17	2.89 ± 0.10	7	52(0.14)	61	0.06	3 2
-0.181	1/4	(2.9; 4.8)	1.09 ± 0.21	3.93 ± 0.13	7	47(0.13)	54	0.07	3 2
-0.215	1/4	(4.3; 7.1)	0.91 ± 0.12	5.59 ± 0.11	7	94(0.11)	109	0.05	4 1
-0.227	1/4	(4.4; 8.2)	0.96 ± 0.10	6.31 ± 0.11	7	116(0.11)	130	0.08	3 2
-0.238	1/4	(5.2; 8.6)	0.88 ± 0.11	6.86 ± 0.12	6	107(0.12)	110	0.25	3 1
-0.249	1/4	(6.1; 9.1)	0.93 ± 0.17	7.60 ± 0.16	6	86(0.12)	98	0.08	3 1
-0.261	1/4	(6.7; 10.1)	1.03 ± 0.19	8.32 ± 0.20	6	70(0.12)	70	0.28	3 1
-0.272	1/4	(7.2; 10.8)	0.84 ± 0.13	9.09 ± 0.22	7	59(0.13)	65	0.11	3 2
-0.283	1/4	(7.5; 11.3)	0.90 ± 0.15	9.36 ± 0.20	6	55(0.14)	63	0.08	3 1
-0.295	1/4	(7.7; 11.5)	0.96 ± 0.16	9.58 ± 0.21	6	54(0.14)	39	0.81	3 1
-0.306	1/4	(8.2; 12.2)	0.90 ± 0.17	10.19 ± 0.22	6	49(0.14)	38	0.69	3 1
-0.317	1/4	(8.0; 13.4)	1.07 ± 0.17	10.72 ± 0.26	6	60(0.14)	59	0.28	3 1
-0.329	1/4	(9.0; 13.4)	1.10 ± 0.24	11.17 ± 0.30	6	45(0.14)	33	0.74	3 1

Table A.4: Fit results for $L \leq 22$ and $q = 1$.

E_i	λ	$(n_0; n_f)$	ν	n_c	N_p	N_D	χ^2	p	$m_L m_\rho$
0.000	1/2	(8.2; 13.6)	0.52 ± 0.19	10.94 ± 0.49	6	59(0.11)	51	0.57	3 1
-0.011	1/2	(8.1; 13.5)	0.49 ± 0.13	10.82 ± 0.39	6	60(0.12)	51	0.61	3 1
-0.023	1/2	(8.8; 13.2)	0.57 ± 0.22	10.99 ± 0.44	6	44(0.12)	48	0.13	3 1
-0.034	1/2	(8.3; 12.4)	0.54 ± 0.18	10.34 ± 0.35	6	48(0.13)	43	0.43	3 1
-0.045	1/2	(7.7; 11.5)	0.80 ± 0.28	9.61 ± 0.38	6	54(0.15)	60	0.12	3 1
-0.057	1/2	(6.8; 10.2)	0.71 ± 0.23	8.56 ± 0.34	6	68(0.15)	80	0.06	3 1
-0.068	1/2	(5.5; 10.1)	0.72 ± 0.13	7.78 ± 0.24	6	117(0.15)	129	0.11	3 1
-0.079	1/2	(4.6; 8.5)	0.98 ± 0.26	6.44 ± 0.27	6	119(0.14)	136	0.07	3 1
-0.091	1/2	(0.0; 1.0)	1.13 ± 0.79	0.48 ± 0.14	6	27(0.62)	32	0.06	3 1
-0.102	1/2	(0.1; 1.3)	0.89 ± 0.30	0.68 ± 0.08	6	37(0.67)	18	0.97	3 1
-0.113	1/2	(0.1; 1.3)	0.74 ± 0.22	0.73 ± 0.07	6	37(0.96)	34	0.31	3 1
-0.125	1/2	(0.4; 1.4)	0.76 ± 0.22	0.89 ± 0.08	6	45(1.21)	40	0.44	3 1
-0.136	1/2	(0.6; 1.7)	0.85 ± 0.28	1.11 ± 0.09	6	48(1.09)	50	0.19	3 1
-0.147	1/2	(0.7; 2.0)	1.49 ± 0.54	1.36 ± 0.13	6	51(1.15)	55	0.15	3 1
-0.159	1/2	(0.9; 2.6)	1.23 ± 0.32	1.69 ± 0.11	6	56(0.87)	60	0.16	3 1
-0.170	1/2	(1.1; 3.2)	1.22 ± 0.27	2.10 ± 0.12	6	62(0.79)	63	0.24	3 1
-0.181	1/2	(1.3; 3.8)	1.48 ± 0.32	2.46 ± 0.14	6	68(0.71)	74	0.14	3 1
-0.193	1/2	(1.6; 4.7)	1.49 ± 0.28	3.07 ± 0.16	6	75(0.54)	88	0.06	3 1
-0.204	1/2	(2.1; 4.9)	1.59 ± 0.37	3.52 ± 0.19	6	66(0.45)	78	0.06	3 1

Table A.4 continued from previous page

E_i	λ	$(n_0; n_f)$	ν	n_c	N_p	N_D (prec.)	χ^2	p	$m_L m_\rho$
-0.215	1/2	(3.0; 5.6)	1.58 ± 0.48	4.32 ± 0.23	6	64(0.34)	57	0.52	3 1
-0.227	1/2	(3.5; 6.5)	1.55 ± 0.40	4.98 ± 0.23	6	87(0.29)	97	0.10	3 1
-0.238	1/2	(4.1; 7.5)	1.72 ± 0.39	5.81 ± 0.24	6	112(0.26)	106	0.48	3 1
-0.249	1/2	(5.0; 8.4)	1.28 ± 0.33	6.70 ± 0.25	6	106(0.24)	96	0.59	3 1
-0.261	1/2	(6.2; 9.2)	1.30 ± 0.45	7.66 ± 0.33	6	83(0.23)	86	0.22	3 1
-0.272	1/2	(6.3; 10.5)	1.58 ± 0.53	8.43 ± 0.47	6	91(0.24)	74	0.81	3 1
-0.283	1/2	(6.1; 11.3)	1.52 ± 0.35	8.61 ± 0.40	6	107(0.27)	104	0.41	3 1
-0.295	1/2	(6.6; 11.0)	1.39 ± 0.40	8.81 ± 0.42	6	84(0.29)	73	0.63	3 1
-0.317	1/2	(8.0; 12.0)	1.24 ± 0.49	9.95 ± 0.49	6	52(0.28)	50	0.32	3 1
0.000	1/4	(7.3; 13.5)	0.45 ± 0.07	10.44 ± 0.26	6	79(0.12)	72	0.51	3 1
-0.011	1/4	(7.4; 13.7)	0.51 ± 0.07	10.45 ± 0.23	6	77(0.12)	77	0.29	3 1
-0.023	1/4	(7.4; 13.7)	0.69 ± 0.09	10.49 ± 0.25	6	77(0.13)	80	0.23	3 1
-0.034	1/4	(6.9; 12.7)	0.68 ± 0.08	9.70 ± 0.19	6	88(0.13)	102	0.07	3 1
-0.045	1/4	(7.7; 10.4)	0.72 ± 0.24	9.14 ± 0.25	6	42(0.13)	40	0.29	3 1
-0.057	1/4	(6.4; 9.6)	0.66 ± 0.11	7.98 ± 0.15	6	79(0.13)	72	0.49	3 1
-0.068	1/4	(5.4; 8)	0.82 ± 0.18	6.74 ± 0.15	6	92(0.13)	96	0.21	3 1
-0.079	1/4	(2.9; 6.1)	1.55 ± 0.33	4.50 ± 0.21	6	83(0.15)	92	0.12	3 1
-0.102	1/4	(0.5; 1.9)	1.24 ± 0.22	1.09 ± 0.08	7	57(0.32)	54	0.33	4 1
-0.125	1/4	(0.4; 2.0)	0.99 ± 0.13	1.21 ± 0.06	6	63(0.53)	72	0.09	3 1
-0.136	1/4	(0.6; 2.2)	1.15 ± 0.21	1.40 ± 0.07	6	61(0.56)	64	0.20	3 1
-0.147	1/4	(1.2; 2.2)	1.48 ± 0.65	1.74 ± 0.13	6	34(0.37)	37	0.13	3 1
-0.170	1/4	(1.6; 3.4)	1.40 ± 0.31	2.52 ± 0.13	6	44(0.32)	39	0.43	3 1
-0.181	1/4	(2.0; 4.1)	1.65 ± 0.39	3.00 ± 0.16	6	51(0.29)	46	0.45	3 1
-0.193	1/4	(2.5; 5.1)	1.37 ± 0.21	3.86 ± 0.13	7	63(0.24)	70	0.10	3 2
-0.204	1/4	(3.1; 5.7)	1.11 ± 0.17	4.29 ± 0.12	6	67(0.22)	77	0.08	3 1
-0.215	1/4	(3.5; 6.5)	1.22 ± 0.19	4.94 ± 0.14	6	87(0.20)	73	0.73	3 1
-0.227	1/4	(4.5; 6.7)	1.14 ± 0.28	5.73 ± 0.16	6	74(0.19)	81	0.13	3 1
-0.238	1/4	(4.8; 8.0)	1.02 ± 0.15	6.33 ± 0.13	6	104(0.19)	95	0.57	3 1
-0.249	1/4	(5.7; 8.5)	1.03 ± 0.22	7.09 ± 0.17	6	92(0.19)	86	0.49	3 1
-0.261	1/4	(5.9; 9.9)	1.17 ± 0.20	7.85 ± 0.21	6	100(0.20)	94	0.47	3 1
-0.272	1/4	(6.4; 10.6)	1.25 ± 0.22	8.48 ± 0.24	6	90(0.21)	91	0.28	3 1
-0.283	1/4	(6.7; 11.1)	1.20 ± 0.19	8.89 ± 0.23	6	83(0.22)	85	0.25	3 1
-0.295	1/4	(6.8; 11.4)	1.13 ± 0.18	9.04 ± 0.24	6	78(0.23)	51	0.97	3 1
-0.306	1/4	(7.3; 12.1)	1.22 ± 0.20	9.68 ± 0.26	6	68(0.24)	65	0.38	3 1
-0.317	1/4	(8.2; 12.4)	1.11 ± 0.30	10.09 ± 0.29	6	47(0.23)	50	0.15	3 1

Table A.5: Fit results for $L \leq 22$ and $q = -1$.

E_i	λ	$(n_0; n_f)$	ν	n_c	N_p	N_D	χ^2	p	$m_L m_\rho$
0.000	1/2	(7.6; 14.0)	0.53 ± 0.13	10.81 ± 0.43	6	74(0.31)	65	0.58	3 1
-0.011	1/2	(7.8; 14.6)	0.62 ± 0.15	11.17 ± 0.45	6	72(0.31)	68	0.40	3 1
-0.023	1/2	(8.3; 13.8)	0.67 ± 0.18	11.01 ± 0.42	6	58(0.31)	67	0.08	3 1

Table A.5 continued from previous page

E_i	λ	$(n_0; n_f)$	ν	n_c	N_p	N_D (prec.)	χ^2	p	m_L m_ρ
-0.034	1/2	(7.8; 13.0)	0.67 ± 0.16	10.44 ± 0.36	6	62(0.33)	64	0.22	3 1
-0.045	1/2	(7.4; 12.3)	0.74 ± 0.17	9.73 ± 0.34	6	66(0.34)	69	0.19	3 1
-0.057	1/2	(6.4; 10.8)	0.70 ± 0.14	8.66 ± 0.28	6	89(0.34)	103	0.07	3 1
-0.068	1/2	(5.8; 9.6)	0.81 ± 0.18	7.67 ± 0.26	6	104(0.34)	112	0.16	3 1
-0.079	1/2	(4.7; 7.9)	1.10 ± 0.39	6.30 ± 0.29	6	103(0.34)	105	0.28	3 1
-0.091	1/2	(1.6; 4.7)	1.22 ± 0.30	3.05 ± 0.23	6	75(0.49)	75	0.29	3 1
-0.102	1/2	(0.8; 2.4)	1.37 ± 0.39	1.65 ± 0.15	6	55(0.78)	51	0.40	3 1
-0.113	1/2	(0.7; 2.0)	1.05 ± 0.27	1.24 ± 0.10	6	51(0.93)	54	0.17	3 1
-0.136	1/2	(1.0; 2.2)	0.74 ± 0.19	1.62 ± 0.10	6	46(1.07)	33	0.78	3 1
-0.147	1/2	(1.2; 2.8)	0.88 ± 0.19	1.95 ± 0.11	6	48(0.93)	53	0.11	3 1
-0.159	1/2	(1.4; 3.4)	0.88 ± 0.18	2.36 ± 0.12	6	51(0.93)	48	0.33	3 1
-0.170	1/2	(1.8; 3.6)	1.02 ± 0.26	2.69 ± 0.14	6	48(0.88)	41	0.49	3 1
-0.181	1/2	(2.0; 4.2)	0.88 ± 0.18	3.11 ± 0.13	6	52(0.87)	46	0.48	3 1
-0.193	1/2	(2.6; 4.8)	1.08 ± 0.28	3.69 ± 0.18	6	55(0.77)	57	0.21	3 1
-0.204	1/2	(3.0; 5.0)	0.78 ± 0.21	4.08 ± 0.15	6	48(0.76)	35	0.77	3 1
-0.215	1/2	(3.4; 6.4)	1.18 ± 0.27	4.86 ± 0.20	6	82(0.61)	75	0.50	3 1
-0.227	1/2	(3.5; 7.3)	1.09 ± 0.15	5.39 ± 0.15	6	117(0.60)	129	0.11	3 1
-0.238	1/2	(4.7; 7.9)	0.97 ± 0.20	6.26 ± 0.17	6	103(0.53)	74	0.96	3 1
-0.249	1/2	(5.2; 8.6)	1.15 ± 0.28	6.83 ± 0.23	6	107(0.56)	99	0.53	3 1
-0.261	1/2	(5.8; 9.6)	1.21 ± 0.30	7.71 ± 0.31	6	104(0.56)	107	0.26	3 1
-0.272	1/2	(5.9; 11.1)	1.41 ± 0.30	8.54 ± 0.37	6	112(0.57)	105	0.51	3 1
-0.283	1/2	(6.1; 11.3)	1.25 ± 0.24	8.67 ± 0.34	6	107(0.60)	85	0.87	3 1
-0.295	1/2	(6.3; 11.7)	1.26 ± 0.25	8.90 ± 0.34	6	101(0.62)	96	0.44	3 1
-0.317	1/2	(7.1; 13.3)	1.46 ± 0.36	10.20 ± 0.52	6	81(0.65)	77	0.41	3 1
-0.329	1/2	(7.9; 13.3)	1.38 ± 0.48	10.50 ± 0.54	6	60(0.67)	39	0.94	3 1

Bibliography

- Abrahams, E., Anderson, P. W., Licciardello, D. C. and Ramakrishnan, T. V. (1979), ‘Scaling theory of localization: Absence of quantum diffusion in two dimensions’, *Physical Review Letters* **42**(10), 673–676. (Cited on pages 3, 67, and 68.)
- Abrahams, E., Kravchenko, S. V. and Sarachik, M. P. (2001), ‘Metallic behavior and related phenomena in two dimensions’, *Reviews of Modern Physics* **73**(April), 251. (Cited on page 10.)
- Abramowitz, M. and Stegun, I. A. (1970), *Handbook of mathematical functions : with formulas, graphs, and mathematical tables*, Dover Publications, New York :. (Cited on page 76.)
- Alexander, M. N. and Holcomb, D. F. (1968), ‘Semiconductor-to-Metal Transition in n-Type Group IV Semiconductors’, *Reviews of Modern Physics* **40**(4), 815–829. (Cited on pages 90 and 94.)
- Al’tshuler, B. L. and Lee, P. A. (1988), ‘Disordered Electronic Systems’, *Physics Today* **41**(12), 36–44. (Cited on page 2.)
- Amini, M., Kravtsov, V. E. and Müller, M. (2014), ‘Multifractality and quantum-to-classical crossover in the Coulomb anomaly at the Mott–Anderson metal–insulator transition’, *New Journal of Physics* **16**(1), 015022. (Cited on pages 11 and 100.)
- Andersen, O. K. and Jepsen, O. (1984), ‘Explicit, first-principles tight-binding theory’, *Physical Review Letters* **53**(27), 2571–2574. (Cited on page 40.)
- Anderson, P. W. (1958), ‘Absence of Diffusion in Certain Random Lattices’, *Physical Review* **109**(5), 1492–1505. (Cited on pages 1, 3, and 102.)

- Anderson, P. W. (1985), 'The question of classical localization A theory of white paint?', *Philosophical Magazine Part B* **52**(3), 505–509. (Cited on pages 7 and 8.)
- Anderson, P. W. (2010), Thoughts on localization, in '50 Years of Anderson Localization', Vol. 24, World Scientific Publishing Co. Pte. Ltd., pp. 1–6. (Cited on page 2.)
- Aoki, H. (1982), 'Decimation method of real-space renormalization for electron systems with application to random systems', *Physica A: Statistical Mechanics and its Applications* **114**(1–3), 538–542. (Cited on page 52.)
- Aoki, H. (1983), 'Critical behaviour of extended states in disordered systems', *Journal of Physics C: Solid State Physics* **16**(6), L205–L208. (Cited on page 52.)
- Aubry, A., Cobus, L. A., Skipetrov, S. E., Van Tiggelen, B. A., Derode, A. and Page, J. H. (2014), 'Recurrent scattering and memory effect at the Anderson localization transition', *Physical Review Letters* **112**(4), 1–5. (Cited on page 8.)
- Aubry, S. and André, G. (1980), 'Analyticity breaking and Anderson localization in incommensurate lattices', *Annals of the Israel Physical Society* **3**(133), 18. (Cited on page 9.)
- Bachelet, G. B. (1986), Electronic States and Structural Properties of Deep Centers in Semiconductors, in 'Crystalline Semiconducting Materials and Devices', Springer US, Boston, MA, pp. 243–304. (Cited on pages 31 and 34.)
- Basko, D., Aleiner, I. and Altshuler, B. (2006), 'Metal–insulator transition in a weakly interacting many–electron system with localized single–particle states', *Annals of Physics* **321**(5), 1126–1205. (Cited on page 10.)
- Belitz, D. and Kirkpatrick, T. R. (1994), 'The Anderson–Mott transition', *Reviews of Modern Physics* **66**(2), 261–380. (Cited on page 10.)
- Bhatt, R. N. and Rice, T. M. (1981), 'Single-particle energy levels in doped semiconductors at densities below the metal–nonmetal transition', *Physical Review B* **23**(4), 1920–1935. (Cited on page 90.)
- Billy, J., Josse, V., Zuo, Z., Bernard, A., Hambrecht, B., Lugan, P., Clément, D., Sanchez-Palencia, L., Bouyer, P. and Aspect, A. (2008), 'Direct observation of Anderson localization of matter waves in a controlled disorder', *Nature* **453**(7197), 891–894. (Cited on page 9.)

- Bishop, D. J., Tsui, D. C. and Dynes, R. C. (1980), 'Nonmetallic Conduction in Electron Inversion Layers at Low Temperatures', *Physical Review Letters* **44**(17), 1153–1156. (Cited on page 3.)
- Bogdanovich, S., Sarachik, M. P. and Bhatt, R. N. (1999), 'Scaling of the Conductivity with Temperature and Uniaxial Stress in Si:B at the Metal-Insulator Transition', *Physical Review Letters* **82**(1), 137–140. (Cited on page 5.)
- Bollhöfer, M. and Notay, Y. (2007), 'JADAMILU: a software code for computing selected eigenvalues of large sparse symmetric matrices', *Computer Physics Communications* **177**(12), 951–964. (Cited on page 48.)
- Born, M. and Oppenheimer, R. (1927), 'Zur Quantentheorie der Molekeln', *Annalen der Physik* **389**(20), 457–484. (Cited on page 17.)
- Brankov, J. G., Danchev, D. M. and Tonchev, N. S. (2000), *Theory of Critical Phenomena in Finite-Size Systems - Scaling and Quantum Effects*, Vol. 9 of *Singapore Perspectives*, World Scientific Publishing Co. Pte. Ltd. (Cited on page 70.)
- Burmistrov, I. S., Gornyi, I. V. and Mirlin, A. D. (2013), 'Multifractality at Anderson Transitions with Coulomb Interaction', *Physical Review Letters* **111**(6), 066601. (Cited on pages 11, 12, and 100.)
- Cardy, J. L. (1996), *Scaling and renormalization in statistical physics*, Cambridge University Press. (Cited on pages 3, 51, 68, 71, 72, and 80.)
- Carnio, E. G., Hine, N. D. M. and Römer, R. A. (2017), 'Resolution of the exponent puzzle for the Anderson transition in doped semiconductors'. (Cited on pages 41, 43, 60, 73, 78, 85, 92, 93, and 95.)
- Castellani, C. and Peliti, L. (1986), 'Multifractal wavefunction at the localisation threshold', *Journal of Physics A: Mathematical and General* **19**(8), L429. (Cited on pages 52 and 55.)
- Chalker, J. T. and Coddington, P. D. (1988), 'Percolation, quantum tunneling and the integer Hall effect', *J. Phys. C* **21**, 2665. (Cited on page 7.)
- Chayes, J. T., Chayes, L., Fisher, D. S. and Spencer, T. (1986), 'Finite-Size Scaling and Correlation Lengths for Disordered Systems', *Physical Review Letters* **57**(24), 2999–3002. (Cited on pages 6 and 103.)

- Chhabra, A. B., Meneveau, C., Jensen, R. V. and Sreenivasan, K. R. (1989), 'Direct determination of the $f(\alpha)$ singularity spectrum and its application to fully developed turbulence', *Physical Review A: Atomic, Molecular, and Optical Physics* **40**(9), 5284–5294. (Cited on page 57.)
- Chhabra, A. B. and Sreenivasan, K. R. (1991), 'Negative dimensions: Theory, computation, and experiment', *Physical Review A* **43**(2), 1114–1117. (Cited on page 60.)
- Clément, D., Varón, A. F., Retter, J. A., Sanchez-Palencia, L., Aspect, A. and Bouyer, P. (2006), 'Experimental study of the transport of coherent interacting matter-waves in a 1D random potential induced by laser speckle', *New Journal of Physics* **8**. (Cited on page 9.)
- Cobus, L., Skipetrov, S., Aubry, A., van Tiggelen, B., Derode, A. and Page, J. (2016), 'Anderson Mobility Gap Probed by Dynamic Coherent Backscattering', *Physical Review Letters* **116**(19), 193901. (Cited on page 8.)
- Coester, F. and Kümmel, H. (1960), 'Short-range correlations in nuclear wave functions', *Nuclear Physics* **17**, 477–485. (Cited on page 22.)
- Cohen, M. H. and Heine, V. (1961), 'Cancellation of Kinetic and Potential Energy in Atoms, Molecules, and Solids', *Physical Review* **122**(6), 1821–1826. (Cited on page 21.)
- Cohen-Tannoudji, C., Diu, B. and Laloë, F. (2005), *Quantum mechanics*, Wiley. (Cited on page 21.)
- Cuevas, E. and Kravtsov, V. E. (2007), 'Two-eigenfunction correlation in a multifractal metal and insulator', *Physical Review B* **76**(23), 235119. (Cited on pages 54 and 72.)
- Cyrot-Lackmann, F. and Gaspard, J. P. (1974), 'On the electronic structure of an impurity band: a cumulant approach', *Journal of Physics C: Solid State Physics* **7**(10), 1829–1839. (Cited on page 91.)
- Dai, P., Zhang, Y. and Sarachik, M. P. (1991), 'Critical conductivity exponent for Si:B', *Physical Review Letters* **66**(14), 1914–1917. (Cited on page 5.)
- del Rio, R., Jitomirskaya, S., Last, Y. and Simon, B. (1995), 'What is Localization?', *Physical Review Letters* **75**(1), 117–119. (Cited on page 39.)

- Devakul, T. and Huse, D. A. (2017), ‘Anderson localization transitions with and without random potentials’, *Physical Review B* **96**(21), 214201. (Cited on page 79.)
- Diu, B., Guthmann, C., Lederer, D. and Roulet, B. (1989), *Éléments de physique statistique*, Hermann. (Cited on page 37.)
- Dobrosavljević, V., Trivedi, N. and Valles, J. M. (2012), *Conductor-Insulator Quantum Phase Transitions*, Oxford University Press, Oxford. (Cited on page 71.)
- Dolan, G. J. and Osheroff, D. D. (1979), ‘Nonmetallic Conduction in Thin Metal Films at Low Temperatures’, *Physical Review Letters* **43**(10), 721–724. (Cited on page 3.)
- Edwards, J. T. and Thouless, D. J. (1972), ‘Numerical studies of localization in disordered systems’, *Journal of Physics C: Solid State Physics* **5**(8), 807–820. (Cited on page 68.)
- Engel, E. and Dreizler, R. M. (2011), *Density Functional Theory, Theoretical and Mathematical Physics*, Springer Berlin Heidelberg, Berlin, Heidelberg. (Cited on pages 18 and 20.)
- Evers, F. and Mirlin, A. D. (2008), ‘Anderson transitions’, *Reviews of Modern Physics* **80**(4), 1355–1417. (Cited on pages 4, 6, 58, 61, and 100.)
- Faez, S., Strybulevych, A., Page, J. H., Lagendijk, A. and van Tiggelen, B. A. (2009), ‘Observation of Multifractality in Anderson Localization of Ultrasound’, *Physical Review Letters* **103**(15), 155703. (Cited on pages 8 and 56.)
- Feher, G., Fletcher, R. C. and Gere, E. A. (1955), ‘Exchange Effects in Spin Resonance of Impurity Atoms in Silicon’, *Physical Review* **100**(6), 1784–1786. (Cited on page 1.)
- Feng, D. and Jin, G. (2005), *Introduction to Condensed Matter Physics (Volume 1)*, World Scientific. (Cited on pages 1, 16, 17, 18, 21, 26, 29, and 30.)
- Finkel'shtein, A. M. (1983), ‘Effect of Coulomb Interaction on the Properties of Disordered Metals’, *Zhurnal Eksperimentalnoi I Teoreticheskoi Fiziki* **84**(1), 168–189. (Cited on page 10.)
- Finkel'stein, A. M. (2010), ‘Disordered electron liquid with interactions’, *International Journal of Modern Physics B* **24**(12n13), 1855–1894. (Cited on page 11.)

- Fisher, M. and Barber, M. (1972), ‘Scaling Theory for Finite-Size Effects in the Critical Region’, *Phys. Rev. Lett.* **28**(23), 1516–1519. (Cited on page 70.)
- Fleishman, L. and Anderson, P. W. (1980), ‘Interactions and the Anderson transition’, *Physical Review B* **21**(6), 2366–2377. (Cited on page 10.)
- Frisch, U. and Parisi, G. (1985), On the singularity structure of fully developed turbulence, in M. Ghil, R. Benzi and G. Parisi, eds, ‘Turbulence and predictability in geophysical fluid dynamics and climate dynamics, Proceedings of the International School of Physics ‘E. Fermi’ (Varenna, Italy, 1983)’, Amsterdam: North-Holland, pp. 84–87. (Cited on page 55.)
- Fritzsche, H. (1958), ‘Resistivity and hall coefficient of antimony-doped germanium at low temperatures’, *Journal of Physics and Chemistry of Solids* **6**(1), 69–80. (Cited on page 89.)
- Ghazali, A. and Hugon, P. L. (1978), ‘Density-Functional Approach to the Metal-Insulator Transition in Doped Semiconductors’, *Physical Review Letters* **41**(22), 1569–1572. (Cited on page 11.)
- Giustino, F. (2014), *Materials Modelling using Density Functional Theory: Properties and Predictions*, Oxford University Press. (Cited on page 20.)
- Görling, A. (1996), ‘Density-functional theory for excited states’, *Physical Review A* **54**(5), 3912–3915. (Cited on page 20.)
- Gornyi, I. V., Mirlin, A. D. and Polyakov, D. G. (2005), ‘Interacting electrons in disordered wires: Anderson localization and low-T transport’, *Physical Review Letters* **95**(20), 1–4. (Cited on page 10.)
- Green, M. A. (1990), ‘Intrinsic concentration, effective densities of states, and effective mass in silicon’, *Journal of Applied Physics* **67**(6), 2944–2954. (Cited on page 26.)
- Greenwood, D. A. (1958), ‘The Boltzmann Equation in the Theory of Electrical Conduction in Metals’, *Proceedings of the Physical Society* **71**(4), 585–596. (Cited on page 68.)
- Harashima, Y. and Slevin, K. (2014), ‘Critical exponent of metal-insulator transition in doped semiconductors: The relevance of the Coulomb interaction’,

Physical Review B **89**, 205108. (Cited on pages 11, 12, 48, 97, 100, 101, 102, and 103.)

Harris, A. B. (1974), ‘Effect of random defects on the critical behaviour of Ising models’, *Journal of Physics C: Solid State Physics* **7**(9), 1671–1692. (Cited on page 6.)

Haynes, P. D., Mostofi, A. A., Skylaris, C.-K. and Payne, M. C. (2006), ‘ONETEP: linear-scaling density-functional theory with plane-waves’, *Journal of Physics: Conference Series* **26**, 143–148. (Cited on page 22.)

Hellmann, H. (1935), ‘A New Approximation Method in the Problem of Many Electrons’, *The Journal of Chemical Physics* **3**(1), 61–61. (Cited on page 17.)

Hellmann, H. and Kassatotschkin, W. (1936), ‘Metallic Binding According to the Combined Approximation Procedure’, *Journal of Chemical Physics* **4**(5), 324–325. (Cited on page 17.)

Herbut, I. (2007), *A Modern Approach to Critical Phenomena*, Cambridge University Press, Cambridge. (Cited on pages 68 and 70.)

Hildebrand, W. K., Strybulevych, A., Skipetrov, S. E., Van Tiggelen, B. A. and Page, J. H. (2014), ‘Observation of infinite-range intensity correlations above, at, and below the mobility edges of the 3D Anderson localization transition’, *Physical Review Letters* **112**(7), 1–5. (Cited on page 8.)

Hirsch, M. J. and Holcomb, D. F. (1987), The Metal-Insulator Transition in Compensated Silicon, in ‘Disordered Semiconductors’, Springer US, Boston, MA, pp. 45–56. (Cited on page 90.)

Hohenberg, P. and Kohn, W. (1964), ‘Inhomogeneous Electron Gas’, *Physical Review* **136**(3B), B864–B871. (Cited on page 18.)

Hu, H., Strybulevych, A., Page, J. H., Skipetrov, S. E. and van Tiggelen, B. A. (2008), ‘Localization of ultrasound in a three-dimensional elastic network’, *Nature Physics* **4**(12), 945–948. (Cited on page 8.)

Huckestein, B. (1995), ‘Scaling theory of the integer quantum Hall effect’, *Reviews of Modern Physics* **67**(2), 357–396. (Cited on page 4.)

- Hung, C. S. and Gliessman, J. R. (1950), ‘The Resistivity and Hall Effect of Germanium at Low Temperatures’, *Physical Review* **79**(4), 726–727. (Cited on page 89.)
- Imada, M., Fujimori, A. and Tokura, Y. (1998), ‘Metal-insulator transitions’, *Reviews of Modern Physics* **70**(4), 1039–1263. (Cited on page 68.)
- Ioffe, A. F. and Regel, A. R. (1960), ‘Non-crystalline, amorphous and liquid electronic semiconductors’, *Progress in Semiconductors* **4**(237), 89. (Cited on page 1.)
- Janssen, M. (1994), ‘Multifractal Analysis of Broadly-Distributed Observables At Criticality’, *International Journal of Modern Physics B* **08**(08), 943–984. (Cited on pages 51, 55, and 61.)
- Jaros, M. (1980), ‘Deep levels in semiconductors’, *Advances in Physics* **29**(3), 409–525. (Cited on page 31.)
- Jendrzejewski, F., Bernard, A., Müller, K., Cheinet, P., Josse, V., Piraud, M., Pezzé, L., Sanchez-Palencia, L., Aspect, A. and Bouyer, P. (2012), ‘Three-dimensional localization of ultracold atoms in an optical disordered potential’, *Nature Physics* **8**(5), 398–403. (Cited on page 9.)
- John, S. (1984), ‘Electromagnetic Absorption in a Disordered Medium near a Photon Mobility Edge’, *Physical Review Letters* **53**(22), 2169–2172. (Cited on pages 7 and 8.)
- John, S., Sompolinsky, H. and Stephen, M. J. (1983), ‘Localization in a disordered elastic medium near two dimensions’, *Physical Review B* **27**(9), 5592–5603. (Cited on page 7.)
- Johnson, T. H., Clark, S. R. and Jaksch, D. (2014), ‘What is a quantum simulator?’, *EPJ Quantum Technology* **1**(10). (Cited on page 10.)
- Jung, D., Czycholl, G. and Kettemann, S. (2012), ‘Finite Size Scaling of the Typical Density of States of Disordered Systems Within the Kernel Polynomial Method’, *International Journal of Modern Physics: Conference Series* **11**, 108–113. (Cited on page 105.)
- Jung, D., Kettemann, S. and Slevin, K. (2016), ‘Anderson metal-insulator transitions with classical magnetic impurities’, *Physical Review B* **93**(13), 134203. (Cited on page 105.)

- Kittel, C. (2005), *Introduction to Solid State Physics*, 8th edn, John Wiley & Sons, Inc. (Cited on page 37.)
- Kohn, W. (1964), ‘Theory of the insulating state’, *Physical Review* **133**(1A). (Cited on page 2.)
- Kohn, W. (1996), ‘Density functional and density matrix method scaling linearly with the number of atoms’, *Physical Review Letters* **76**(17), 3168–3171. (Cited on page 22.)
- Kohn, W. and Sham, L. J. (1965), ‘Self-consistent equations including exchange and correlation effects’, *Physical Review* **140**(4A). (Cited on pages 11 and 18.)
- Kondov, S. S., McGehee, W. R., Zirbel, J. J. and DeMarco, B. (2011), ‘Three-Dimensional Anderson Localization of Ultracold Matter’, *Science* **334**(6052), 66–68. (Cited on page 9.)
- Kramer, B., Ohtsuki, T. and Kettemann, S. (2005), ‘Random network models and quantum phase transitions in two dimensions’, *Physics Reports* **417**(5–6), 211–342. (Cited on page 4.)
- Kravchenko, S. V., Mason, W. E., Bowker, G. E., Furneaux, J. E., Pudalov, V. M. and D’Iorio, M. (1995), ‘Scaling of an anomalous metal-insulator transition in a two-dimensional system in silicon at $B=0$ ’, *Physical Review B* **51**(11), 7038–7045. (Cited on page 4.)
- Kreutzer, M., Pieper, A., Hager, G., Wellein, G., Alvermann, A. and Fehske, H. (2015), ‘Performance Engineering of the Kernel Polynomial Method on Large-Scale CPU-GPU Systems’, *Proceedings - 2015 IEEE 29th International Parallel and Distributed Processing Symposium, IPDPS 2015* pp. 417–426. (Cited on page 105.)
- Kubo, R. (1957), ‘Statistical-Mechanical Theory of Irreversible Processes. I. General Theory and Simple Applications to Magnetic and Conduction Problems’, *Journal of the Physical Society of Japan* **12**(6), 570–586. (Cited on page 67.)
- Lagendijk, A., van Tiggelen, B. and Wiersma, D. S. (2009), ‘Fifty years of Anderson localization’, *Physics Today* **62**(8), 24–29. (Cited on page 2.)
- Landau, D. P. and Binder, K. (2000), *A Guide to Monte Carlo Simulations in Statistical Physics*, Cambridge University Press, Cambridge. (Cited on page 34.)

- Landauer, R. (1957), ‘Spatial Variation of Currents and Fields Due to Localized Scatterers in Metallic Conduction’, *IBM Journal of Research and Development* 1(3), 223–231. (Cited on page 3.)
- Landauer, R. (1970), ‘Electrical resistance of disordered one-dimensional lattices’, *Philosophical Magazine* 21(172), 863–867. (Cited on pages 3 and 67.)
- Li, W., Vicente, C. L., Xia, J. S., Pan, W., Tsui, D. C., Pfeiffer, L. N. and West, K. W. (2009), ‘Scaling in plateau-to-plateau transition: A direct connection of quantum hall systems with the Anderson localization model’, *Physical Review Letters* 102(21), 1–4. (Cited on page 7.)
- Li, X. P., Nunes, W. and Vanderbilt, D. (1993), ‘Density-Matrix Electronic-Structure Method With Linear System- Size Scaling’, *Physical Review B* 47(16), 10891–10894. (Cited on page 22.)
- Lindinger, J. and Rodríguez, A. (2017), ‘Multifractal finite-size scaling at the Anderson transition in the unitary symmetry class’, *Physical Review B* 96(13), 134202. (Cited on pages 48, 74, and 80.)
- Lopez, M., Clément, J.-F., Szriftgiser, P., Garreau, J. C. and Delande, D. (2012), ‘Experimental Test of Universality of the Anderson Transition’, *Physical Review Letters* 108(9), 095701. (Cited on page 10.)
- Luque, A. and Martí, A. (1997), ‘Increasing the Efficiency of Ideal Solar Cells by Photon Induced Transitions at Intermediate Levels’, *Physical Review Letters* 78(26), 5014–5017. (Cited on page 91.)
- Luque, A., Martí, A., Antolín, E. and Tablero, C. (2006), ‘Intermediate bands versus levels in non-radiative recombination’, *Physica B: Condensed Matter* 382(1–2), 320–327. (Cited on page 91.)
- M. Itoh, K., Watanabe, M., Ootuka, Y., E. Haller, E. and Ohtsuki, T. (2004), ‘Complete Scaling Analysis of the Metal–Insulator Transition in Ge:Ga: Effects of Doping–Compensation and Magnetic Field’, *Journal of the Physical Society of Japan* 73(1), 173–183. (Cited on pages 5, 6, 101, and 104.)
- MacKinnon, A. and Kramer, B. (1981), ‘One-Parameter Scaling of Localization Length and Conductance in Disordered Systems’, *Physical Review Letters* 47(21), 1546–1549. (Cited on pages 3, 6, 70, 74, and 76.)

- Malcai, O., Lidar, D. A., Biham, O. and Avnir, D. (1997), 'Scaling range and cutoffs in empirical fractals', *Physical Review E* **56**(3), 2817–2828. (Cited on page 53.)
- Mandelbrot, B. B. (1972), Possible refinement of the lognormal hypothesis concerning the distribution of energy dissipation in intermittent turbulence, Springer, Berlin, Heidelberg, pp. 333–351. (Cited on page 52.)
- Mandelbrot, B. B. (1974), 'Intermittent turbulence in self-similar cascades: divergence of high moments and dimension of the carrier', *Journal of Fluid Mechanics* **62**(02), 331. (Cited on page 52.)
- Mandelbrot, B. B. (1984), 'Fractals in physics: Squig clusters, diffusions, fractal measures, and the unicity of fractal dimensionality', *Journal of Statistical Physics* **34**(5–6), 895–930. (Cited on page 60.)
- Martin, R. M. (2004), *Electronic Structure: Basic Theory and Practical Methods*, Cambridge University Press, Cambridge. (Cited on pages 16, 18, 21, and 26.)
- Miller, A. and Abrahams, E. (1960), 'Impurity conduction at low concentrations', *Physical Review* **120**(3), 745–755. (Cited on page 10.)
- Mirlin, A. D. and Evers, F. (2000), 'Multifractality and critical fluctuations at the Anderson transition', *Physical Review B* **62**(12), 7920–7933. (Cited on page 61.)
- Mirlin, A. D., Fyodorov, Y. V., Mildenerger, A. and Evers, F. (2006), 'Exact relations between multifractal exponents at the anderson transition', *Physical Review Letters* **97**(4), 1–4. (Cited on pages 56 and 100.)
- Møller, C. and Plesset, M. S. (1934), 'Note on an Approximation Treatment for Many-Electron Systems', *Physical Review* **46**(7), 618–622. (Cited on page 22.)
- Moore, C., Tam, K. M., Zhang, Y. and Jarrell, M. (2017), 'Multifractal Study of Quasiparticle Localization in Disordered Superconductors', (1), 1–7. (Cited on pages 48 and 79.)
- Mott, N. F. (1966), 'The electrical properties of liquid mercury', *Philosophical Magazine* **13**(125), 989–1014. (Cited on page 2.)
- Mott, N. F. (1990), *Metal-Insulator Transitions*, 2nd ed. edn, Taylor & Francis, London. (Cited on pages 89 and 90.)

- Mott, N. F. and Twose, W. D. (1961), ‘The theory of impurity conduction’, *Advances in Physics* **10**(38), 107–163. (Cited on page 3.)
- Nagy, B. and Strand, R. (2009), ‘Neighborhood sequences in the diamond grid-Algorithms with four neighbors’, *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* **5852 LNCS**, 109–121. (Cited on page 24.)
- Ndawana, M. L., Römer, R. A. and Schreiber, M. (2004), ‘The Anderson metal-insulator transition in the presence of scale-free disorder’, *Europhysics Letters (EPL)* **68**(5), 678–684. (Cited on page 102.)
- Neugebauer, J. and Hickel, T. (2013), ‘Density functional theory in materials science’, *Wiley Interdisciplinary Reviews: Computational Molecular Science* **3**(5), 438–448. (Cited on page 19.)
- O’Regan, D. D. (2012), *Optimised Projections for the Ab Initio Simulation of Large and Strongly Correlated Systems*, Springer Berlin Heidelberg, Berlin, Heidelberg. (Cited on page 23.)
- Overhof, H., Scheffler, M. and Weinert, C. M. (1991), ‘Formation energies, electronic structure, and hyperfine fields of chalcogen point defects and defect pairs in silicon’, *Physical Review B* **43**(15), 12494–12506. (Cited on page 31.)
- Owen, F. and Jones, R. (1994), *Statistics*, Pitman. (Cited on page 76.)
- Paalanen, M. A., Rosenbaum, T. F., Thomas, G. A. and Bhatt, R. N. (1982), ‘Stress tuning of the metal-insulator transition at millikelvin temperatures’, *Physical Review Letters* **48**(18), 1284–1287. (Cited on page 5.)
- Paalanen, M. A., Rosenbaum, T. F., Thomas, G. A. and Bhatt, R. N. (1983), ‘Critical Scaling of the Conductance in a Disordered Insulator’, *Physical Review Letters* **51**(20), 1896–1899. (Cited on page 4.)
- Patterson, J. and Bailey, B. (2010), *Solid-State Physics*, Springer Berlin Heidelberg, Berlin, Heidelberg. (Cited on pages 18 and 21.)
- Pichard, J. L. and Sarma, G. (1981), ‘Finite size scaling approach to Anderson localisation’, *Journal of Physics C: Solid State Physics* **14**(6), L127–L132. (Cited on pages 6 and 70.)

- Pietronero, L. (1990), Fractals in Physics, in 'The Superworld III', Springer US, Boston, MA, pp. 175–183. (Cited on page 52.)
- Plyushchay, I. V., Römer, R. A. and Schreiber, M. (2003), 'Three-dimensional Anderson model of localization with binary random potential', *Physical Review B* **68**(6), 064201. (Cited on page 103.)
- Porezag, D., Frauenheim, T., Köhler, T., Seifert, G. and Kaschner, R. (1995), 'Construction of tight-binding-like potentials on the basis of density-functional theory: Application to carbon', *Physical Review B* **51**(19), 12947–12957. (Cited on page 40.)
- Press, W. H., Teukolsky, S. A., Vetterling, W. T. and Flannery, B. P. (2007), *Numerical recipes : the art of scientific computing*, Cambridge University Press. (Cited on pages 23 and 76.)
- Punnoose, A. and Finkel'stein, A. M. (2005), 'Metal-Insulator Transition in Disordered Two-Dimensional Electron Systems', *Science* **310**(5746), 289–291. (Cited on pages 10 and 104.)
- Puschmann, M., Cain, P., Schreiber, M. and Vojta, T. (2015), 'Multifractal analysis of electronic states on random Voronoi-Delaunay lattices', *European Physical Journal B* **88**(11), 1–9. (Cited on pages 48, 79, and 102.)
- Radjenovic, B. and Tjapkin, D. (1989), 'The Influence of Hybridization of the Impurity and the Conduction Bands on the Density of States of Doped Semiconductors', *physica status solidi (b)* **156**(2), 487–495. (Cited on page 103.)
- Resta, R. (1986), Shallow Impurity States in Semiconductors: Theoretical Aspects, in 'Crystalline Semiconducting Materials and Devices', Springer US, Boston, MA, pp. 217–241. (Cited on page 29.)
- Richardella, A., Roushan, P., Mack, S., Zhou, B., Huse, D. A., Awschalom, D. D. and Yazdani, A. (2010), 'Visualizing Critical Correlations Near the Metal-Insulator Transition in $\text{Ga}_{1-x}\text{Mn}_x\text{As}$ ', *Science* **327**(5966), 665–669. (Cited on pages 11 and 100.)
- Roati, G., D'Errico, C., Fallani, L., Fattori, M., Fort, C., Zaccanti, M., Modugno, G., Modugno, M. and Inguscio, M. (2008), 'Anderson localization of a non-interacting Bose-Einstein condensate', *Nature* **453**(7197), 895–898. (Cited on page 9.)

- Röder, H., Silver, R. N., Drabold, D. A. and Dong, J. J. (1997), ‘Kernel polynomial method for a nonorthogonal electronic-structure calculation of amorphous diamond’, *Physical Review B* **55**(23), 15382–15385. (Cited on page 105.)
- Rodriguez, A., Vasquez, L. J. and Römer, R. a. (2008), ‘Multifractal analysis of the metal-insulator transition in the three-dimensional Anderson model. II. Symmetry relation under ensemble averaging’, *Physical Review B* **78**(19), 195107. (Cited on pages 56, 60, 61, 63, and 100.)
- Rodriguez, A., Vasquez, L. J. and Römer, R. A. (2009), ‘Optimisation of multifractal analysis at the 3D Anderson transition using box-size scaling’, *The European Physical Journal B* **67**(1), 77–82. (Cited on page 48.)
- Rodriguez, A., Vasquez, L. J., Slevin, K. and Römer, R. A. (2011), ‘Multifractal finite-size scaling and universality at the Anderson transition’, *Physical Review B* **84**(13), 134209. (Cited on pages 7, 12, 48, 58, 71, 74, 76, 79, 80, 81, 83, and 103.)
- Rosenbaum, T. F., Andres, K., Thomas, G. A. and Bhatt, R. N. (1980), ‘Sharp Metal-Insulator Transition in a Random Solid’, *Physical Review Letters* **45**(21), 1723–1726. (Cited on page 5.)
- Saad, Y. (2011), *Numerical Methods for Large Eigenvalue Problems*, Society for Industrial and Applied Mathematics, Philadelphia. (Cited on page 47.)
- Sankey, O. F. and Niklewski, D. J. (1989), ‘Ab initio multicenter tight-binding model for molecular-dynamics simulations and other applications in covalent systems’, *Physical Review B* **40**(6), 3979–3995. (Cited on page 40.)
- Scheffold, F., Lenke, R., Tweer, R. and Maret, G. (1999), ‘Localization or classical diffusion of light?’, *Nature* **398**(6724), 206–207. (Cited on page 8.)
- Schenk, O., Bollhöfer, M. and Römer, R. A. (2006), ‘On Large-Scale Diagonalization Techniques for the Anderson Model of Localization’, *SIAM Journal on Scientific Computing* **28**(3), 963–983. (Cited on page 48.)
- Schreiber, M., Kramer, B. and MacKinnon, A. (1989), ‘Localization of Disordered Electron Systems – A Numerical Analysis of the Anderson Transition’, *Physica Scripta* **T25**, 67–71. (Cited on page 6.)

- Schwartz, T., Bartal, G., Fishman, S. and Segev, M. (2007), ‘Transport and Anderson localization in disordered two-dimensional photonic lattices’, *Nature* **446**(7131), 52–55. (Cited on page 8.)
- Segev, M., Silberberg, Y. and Christodoulides, D. N. (2013), ‘Anderson localization of light’, *Nature Photonics* **7**(3), 197–204. (Cited on page 7.)
- Semeghini, G., Landini, M., Castilho, P., Roy, S., Spagnolli, G., Trenkwalder, A., Fattori, M., Inguscio, M. and Modugno, G. (2015), ‘Measurement of the mobility edge for 3D Anderson localization’, *Nature Physics* **11**(7), 554–559. (Cited on pages 9 and 10.)
- Sheng, P. (2006), *Introduction to Wave Scattering, Localisation, and Mesoscopic Phenomena*. (Cited on page 1.)
- Skipetrov, S. E. and Page, J. H. (2016), ‘Red light for Anderson localization’, *New Journal of Physics* **18**(2), 021001. (Cited on page 8.)
- Skylaris, C.-K. and Haynes, P. D. (2007), ‘Achieving plane wave accuracy in linear-scaling density functional theory applied to periodic systems: a case study on crystalline silicon.’, *The Journal of chemical physics* **127**(16), 164712. (Cited on page 24.)
- Slevin, K. and Ohtsuki, T. (1999), ‘Corrections to Scaling at the Anderson Transition’, *Physical Review Letters* **82**(2), 382–385. (Cited on pages 6, 12, 39, 79, and 97.)
- Slevin, K. and Ohtsuki, T. (2009), ‘Critical exponent for the quantum Hall transition’, *Physical Review B - Condensed Matter and Materials Physics* **80**, 1–4. (Cited on page 7.)
- Slevin, K. and Ohtsuki, T. (2014), ‘Critical exponent for the Anderson transition in the three-dimensional orthogonal universality class’, *New Journal of Physics* **16**(1), 015012. (Cited on pages 7 and 79.)
- Smith, J. S., Budi, A., Per, M. C., Vogt, N., Drumm, D. W., Hollenberg, L. C. L., Cole, J. H. and Russo, S. P. (2017), ‘Ab initio calculation of energy levels for phosphorus donors in silicon’, *Scientific Reports* **7**(1), 6010. (Cited on page 30.)
- Sperling, T., Schertel, L., Ackermann, M., Aubry, G. J., Aegerter, C. M. and Maret, G. (2016), ‘Can 3D light localization be reached in ‘white paint?’’, *New Journal of Physics* **18**(1), 013039. (Cited on page 8.)

- Störzer, M., Gross, P., Aegerter, C. M. and Maret, G. (2006), ‘Observation of the Critical Regime Near Anderson Localization of Light’, *Physical Review Letters* **96**(6), 063904. (Cited on page 8.)
- Stupp, H., Hornung, M., Lakner, M., Madel, O. and Löhneysen, H. v. (1993), ‘Possible solution of the conductivity exponent puzzle for the metal-insulator transition in heavily doped uncompensated semiconductors’, *Physical Review Letters* **71**(16), 2634–2637. (Cited on page 5.)
- Tél, T. (1988), ‘Fractals, Multifractals, and Thermodynamics’, *Zeitschrift für Naturforschung A* **43**(12), 1154–1174. (Cited on page 51.)
- Thiem, S. and Schreiber, M. (2013), ‘Partitioning schemes and non-integer box sizes for the box-counting algorithm in multifractal analysis’, *The European Physical Journal B* **86**(2), 48. (Cited on pages 48 and 57.)
- Thomas, G. A. (1985), ‘Critical exponent of the metal-insulator transition’, *Philosophical Magazine Part B* **52**(3), 479–498. (Cited on page 5.)
- Thouless, D. (1974), ‘Electrons in disordered systems and the theory of localization’, *Physics Reports* **13**(3), 93–142. (Cited on pages 2, 67, and 68.)
- Thouless, D. (2010), Anderson localization in the Seventies and beyond, in ‘50 Years of Anderson Localization’, WORLD SCIENTIFIC, pp. 7–25. (Cited on page 11.)
- Uren, M. J., Davies, R. A., Kaveh, M. and Pepper, M. (1981), ‘Logarithmic corrections to two-dimensional transport in silicon inversion layers’, *Journal of Physics C: Solid State Physics* **14**(36), 5737–5762. (Cited on page 3.)
- v. Klitzing, K., Dorda, G. and Pepper, M. (1980), ‘New method for high-accuracy determination of the fine-structure constant based on quantized hall resistance’, *Physical Review Letters* **45**(6), 494–497. (Cited on page 3.)
- Van Hove, L. (1949), ‘Quelques propriétés générales de l’intégrale de configuration d’un système de particules avec interaction’, *Physica* **XV**(11), 951–961. (Cited on page 69.)
- Vasquez, L. J., Rodriguez, A. and Römer, R. A. (2008), ‘Multifractal analysis of the metal-insulator transition in the three-dimensional Anderson model. I.

- Symmetry relation under typical averaging', *Physical Review B* **78**(19), 195106. (Cited on page 56.)
- Waffenschmidt, S., Pfeleiderer, C. and Löhneysen, H. v. (1999), 'Critical Behavior of the Conductivity of Si:P at the Metal-Insulator Transition under Uniaxial Stress', *Physical Review Letters* **83**(15), 3005–3008. (Cited on pages 5 and 6.)
- Wegner, F. (1979), 'The mobility edge problem: Continuous symmetry and a conjecture', *Zeitschrift für Physik B Condensed Matter and Quanta* **35**(3), 207–210. (Cited on page 10.)
- Wegner, F. (1980), 'Inverse participation ratio in $2 + \epsilon$ dimensions', *Zeitschrift für Physik B Condensed Matter and Quanta* **36**(3), 209–214. (Cited on page 52.)
- Wegner, F. J. (1976), 'Electrons in disordered systems. Scaling near the mobility edge', *Zeitschrift für Physik B Condensed Matter and Quanta* **25**(4), 327–337. (Cited on page 67.)
- Weinrib, A. and Halperin, B. I. (1983), 'Critical phenomena in systems with long-range-correlated quenched disorder', *Physical Review B* **27**(1), 413–427. (Cited on page 102.)
- Wellens, T. and Jalabert, R. A. (2016), 'Charge and spin diffusion on the metallic side of the metal-insulator transition: A self-consistent approach', *Physical Review B* **94**(14), 144209. (Cited on page 91.)
- Wiersma, D. S. (2013), 'Disordered photonics', *Nature Photonics* **7**(3), 188–196. (Cited on page 7.)
- Wiersma, D. S., Bartolini, P., Lagendijk, A. and Righini, R. (1997), 'Localization of light in a disordered medium', *Nature* **390**(6661), 671–673. (Cited on page 8.)
- Wilson, K. G. (1971), 'Renormalization group and critical phenomena. II. Phase-space cell analysis of critical behavior', *Physical Review B* **4**(9), 3184–3205. (Cited on page 67.)
- Wilson, N. R., Nguyen, P. V., Seyler, K., Rivera, P., Marsden, A. J., Laker, Z. P., Constantinescu, G. C., Kandyba, V., Barinov, A., Hine, N. D., Xu, X. and Cobden, D. H. (2017), 'Determination of band offsets, hybridization, and exciton binding in 2D semiconductor heterostructures', *Science Advances* **3**(2), e1601832. (Cited on page 104.)

- Winkler, M. T., Recht, D., Sher, M.-J., Said, A. J., Mazur, E. and Aziz, M. J. (2011), ‘Insulator-to-Metal Transition in Sulfur-Doped Silicon’, *Physical Review Letters* **106**(17), 178701. (Cited on pages 34, 49, 91, and 93.)
- Wolfram Research Inc. (2015), *Mathematica 10.3*, Wolfram Research, Inc., Champaign, Illinois. (Cited on page 46.)
- Yakubo, K. and Ono, M. (1998), ‘Finite-size scaling of multifractal wave functions: The metal-insulator transition in two-dimensional symplectic systems’, *Phys. Rev. B* **58**(15), 9767–9772. (Cited on page 71.)
- Yu, P. Y. and Cardona, M. (2010), *Fundamentals of Semiconductors*, Vol. 1 of *Graduate Texts in Physics*, Springer Berlin Heidelberg, Berlin, Heidelberg. (Cited on pages 24, 29, 30, 31, and 34.)
- Zhang, S., Yamagiwa, S., Okumura, M. and Yunoki, S. (2013), ‘Kernel polynomial method on GPU’, *International Journal of Parallel Programming* **41**(1), 59–88. (Cited on page 105.)