Fakultät für Physik
Institut für Theoretische Festkörperphysik

Band Structure Based Analysis of Certain Photonic Crystal Structures

Dissertation thesis

by

Dipl.–Phys. Christian Wolff

hand–in date: 23.11.2011

Instructor:	Prof. Dr. Kurt Busch
2nd Instructor:	Prof. Dr. Martin Wegener

Band Structure Based Analysis of Certain Photonic Crystal Structures

Zur Erlangung des akademischen Grades eines

DOKTORS DER NATURWISSENSCHAFTEN

von der Fakultät für Physik des
Karlsruher Instituts für Technologie (KIT)

genehmigte

DISSERTATION

von

Dipl. Phys. Christian Wolff
aus Karlsruhe

Tag der mündlichen Prüfung: 09.12.2011
Referent: Prof. Dr. K. Busch
Korreferent: Prof. Dr. M. Wegener

Ich erkläre hiermit, dass ich die vorliegende Arbeit selbständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel verwendet habe.

Karlsruhe, den 28.01.2012

Introduction

Light is a fundamental part of daily life, not only because we can directly perceive it with our eyes. It is utilized in a vast number of applications throughout engineering and science. Some examples are illumination, communication and information transfer or the production of electric energy, which has gained great public interest in the context of solar cell development. All these applications require appropriate light sources or components for guiding and manipulating radiation.

Until a few decades ago, this has been done mainly by combining macroscopic objects made of materials that can be considered homogeneous on the length scale of light. However, in particular the revolutionary advances in semiconductor technology for the fabrication of integrated electronics provide a toolbox to accurately structure materials on a nanometer scale.

Periodic arrangements of dielectric materials that are structured on the length scale of visible light, so–called photonic crystals, were initially introduced to inhibit spontaneous emission [1] by a complete photonic band gap and facilitate Anderson localization [2] of light. Furthermore, they feature highly dispersive and anisotropic transport which may be exploited for negative refraction [3] or the construction of superprisms [4]. Eigenmodes of point–like perturbations or pile–up effects due to a very low group velocity (slow light) may cause large field enhancements. This increases the coupling between light and classical nonlinear materials or quantum systems such as single atoms or quantum dots. In general, the strong internal diffraction occurring inside a photonic crystal tends to make the physics of light propagation very feature–rich. Their dependence on the geometry of the structure provides great freedom to tailor the behavior of light to a specific purpose.

For instance, it is possible to almost deliberately control the density of states in a given frequency range and, thereby, not only to inhibit but also to enhance spontaneous emission. The anisotropic nature of light propagation can be further used to control the direction in which light is emitted. For example, in the spectral vicinity of van–Hove singularities, spontaneous emission can be expected to be enhanced and confined to few directions. Similar possibilities exist for the control of thermal emission [5, 6] with the restriction that the black–body limit cannot be overcome.

Finally, the possibility to deliberately introduce deviations from the ideal crystal (defects) allows to modify its local properties such that a broad variety of functional devices can be implemented. Point defects may be used as high quality resonators, e. g. , in lasers with low threshold [7]. Wave guides can be formed by placing individual resonators in close proximity, thus forming a line defect. More complex defect arrangements allow to couple and filter light or to have two individual waves interact via a nonlinearity on a length scale of only few microns [8]. Thus, defect structures in photonic crystals offer a promising platform for the development of all–optical integrated circuits.

Numerical simulations serve several purposes in the development of photonic crystal structures and their optimization for a specific task. First, they provide a relatively inexpensive and quick way to estimate whether a structure can serve the desired purpose or not. Second, comparing the measured properties of the experimentally fabricated structure with theoretically expected data allows to rate the quality of the sample and may indicate unexpected fabricational imperfections or side effects of the measurement process. General purpose Maxwell solvers such as the finite–difference time–domain method [9] (FDTD), several finite–element–like simulators (e. g. , the discontinuous Galerkin methods [10]) or the Fourier Modal Method [11, 12] (FMM) are excellent choices for these tasks as they can model the experiment in quite some detail. For instance, the simulated spectra may exhibit strong features due to finite–size effects or certain types of disorder.

Physical insight of an effect is most easily gained by systematically eliminating possible uncertainties. In simulations, this usually means to disable physical processes such as energy loss due to absorption. Ab–initio solvers, however, are not well–suited for investigating the details of the complex propagation and scattering processes of light inside a photonic crystal. For instance, in an FDTD calculation it is very difficult to deduce the impact of an individual eigenmode in an interference process.

The central quantity that describes the behavior of light in matter is the dispersion relation of the propagating modes. In the case of (defect–free) photonic crystals, this is the photonic band structure. It provides information about the density of states, energy transport and the phase velocity. Together with the corresponding set of eigenmodes, the photonic band structure is the key to identify the physical origin of most effects related to wave propagation such as refraction and anisotropic propagation. However, it is ill–suited to predict the behavior of localized defect structures, mainly because the eigenmodes resemble plane waves, i. e. , are infinitely extended. Instead, a basis of confined functions is preferred for this type of problems. Photonic Wannier functions [13] form a set of functions constructed from the photonic crystal's eigenmodes and can be chosen to be well localized. They inherit the full information about the photonic band structure and form a very efficient basis to describe the ideal crystal and defect structures such as integrated optical circuits via a perturbatory approach. It has also proven to exhibit fundamental advantages over the eigenmode basis in the context of disorder in photonic crystals [14, 15].

This thesis is organized as follows:

In the first chapter, I briefly introduce the physical background this thesis is based on. This includes the fundamental approximations which I assumed throughout my work, stating Maxwell's equations as the fundamental equations of motion for the electromagnetic fields and a short introduction to photonic crystals and related defect structures.

As the dispersion relation is the central quantity in the context of light propagation, the second chapter is the natural place to discuss some of the most popular approaches to numerical band structure calculations. Besides the plane wave expansion for conventional and on–shell band structure calculations, I discuss the “MIT Photonic Bands” software package [16] and a linear interpolation approach as an approximate on–shell solver.

The third chapter is dedicated to the photonic Wannier function basis, which is derived from the photonic crystal’s eigenmodes. Here, my emphasis is on efficient ways to create Wannier functions that form a well–suited expansion basis for numerical defect calculations. Furthermore, the problem of absorbing boundaries in photonic crystal wave guides is addressed.

While the first three chapters are rather technical, I apply the band structure–based methods to interesting physical problems in the remainder of my thesis. The forth chapter deals with a birefringence–like polarization conversion effect that occurs in thin films of artificial opals. From a conventional crystal optics point of view, this effect is unexpected due to the cubic nature of the opal structure. This is another example for the feature–rich wave propagation in photonic crystals.

In my fifth chapter, I discuss a novel experimental approach to measuring the dynamics of spontaneous emission in a structured material. Furthermore, I derive the theory to calculate the line width and anomalous Lamb shift of a single magnetic dipole emitter inside a band gap material and show numerical results for these two quantities. The theoretical and experimental data confirm the inhibition of spontaneous emission inside the band gap [1] and the existence of non–Markovian dynamics at the band edges [17].

Finally, I conclude my thesis with a summary, an outlook and two appendices containing lengthy derivations.

Contents

1
Chapter 1.
Fundamentals

This chapter is a very short summary of the fundamental physical framework that is required throughout this thesis. In the first section, I state my main approximations and the basic equations that I use to describe classical electrodynamics in dielectric systems. The second section deals with ideal photonic crystals (PhC) being one special class of dielectric structures and the main topic of my thesis. The key feature of PhCs is their periodicity with respect to the refractive index distribution.

1.1. Electromagnetism

In my thesis, I focus on electromagnetic fields of moderate intensity at optical frequencies. Within this context, moderate intensity refers to fields that are so weak that nonlinear response of materials may be neglected but that are strong enough to be described by classical fields rather than by individual photons. The typical length scale in optics is about 100nm which is roughly three orders of magnitude larger than the atomic length scale. Hence, we may regard materials as homogeneous and describe them by averaged quantities such as susceptibilities. Moreover, non–local effects may be neglected as well because the correlation length of electrons is on the atomic scale. In the optical regime, there are virtually no materials that exhibit relevant magnetic response, thus we can safely assume all relative permeabilities to have the value one. Finally, we restrict ourselves to lossless and dispersion–free systems, which for instance excludes metals. To motivate the former assumption, I should mention that losses degrade the performance of many important PhC–based systems such as high–quality nanocavities [18]. Therefore, materials with as low losses as possible are preferred. Due to the Kramers–Kronig relations, a material with only little absorption over a wide spectral range is approximately non–dispersive in the center of this range. This finally motivates the last assumption.

Putting together all points above, we find that our range of parameters is described by the macroscopic Maxwell equations [19] in SI–units:

$$\begin{aligned}\nabla \times \mathbf{E}(\mathbf{r},t) =& -\partial_t \mathbf{B}(\mathbf{r},t), & (1.1)\\ \nabla \times \mathbf{H}(\mathbf{r},t) =& \partial_t \mathbf{D}(\mathbf{r},t) + \mathbf{j}(\mathbf{r},t), & (1.2)\\ \nabla \cdot \mathbf{D}(\mathbf{r},t) =& \rho(\mathbf{r},t), & (1.3)\\ \nabla \cdot \mathbf{B}(\mathbf{r},t) =& 0, & (1.4)\end{aligned}$$

where ρ and $\mathbf{j}$ refer to the free charge distribution and the free current density, respectively. The material properties enter via the linearized, local and memory–free constitutive relations for the electric field $\mathbf{E}$, the electric displacement field $\mathbf{D}$, the magnetic induction $\mathbf{H}$ and the magnetic flux density $\mathbf{B}$:

$$\begin{aligned}\mathbf{D}(\mathbf{r},t) =& \epsilon_0 \epsilon(\mathbf{r}) \mathbf{E}(\mathbf{r},t), & (1.5)\\ \mathbf{B}(\mathbf{r},t) =& \mu_0 \mathbf{H}(\mathbf{r},t). & (1.6)\end{aligned}$$

The relative dielectric function $\epsilon(\mathbf{r})$ is real–valued as an imaginary part would correspond to dissipation; a magnetic material response was neglected.

It is convenient for the investigation of stationary problems to formulate above equations with respect to frequency rather than to time. To this end, we assume harmonic time dependence of the fields:

$$\Psi(\mathbf{r},t) = \Psi(\mathbf{r},\omega) \exp\big(-\mathrm{i}\omega t\big), \tag{1.7}$$

with the real–valued angular frequency[1] ω. Here, Ψ stands for the charge density or a vector component of any electromagnetic field or of the current density. The temporal ansatz with only one spectral component is sufficient because we only consider linear material response. Applying the temporal derivative yields:

$$\partial_t \Big(\Psi(\mathbf{r},\omega) \exp\big(-\mathrm{i}\omega t\big) \Big) = -\mathrm{i}\omega \Psi(\mathbf{r},\omega) \exp\big(-\mathrm{i}\omega t\big). \tag{1.8}$$

By dividing by the oscillating exponential, we find substitution rules for transforming our equations of motion to frequency–domain:

$$\begin{aligned}\Psi(\mathbf{r},t) \rightarrow& \Psi(\mathbf{r},\omega), & (1.9)\\ \partial_t \rightarrow& -\mathrm{i}\omega. & (1.10)\end{aligned}$$

We finally arrive at the stationary Maxwell equations:

$$\begin{aligned}\nabla \times \mathbf{E}(\mathbf{r},\omega) =& \mathrm{i}\omega \mathbf{B}(\mathbf{r},\omega), & (1.11)\\ \nabla \times \mathbf{H}(\mathbf{r},\omega) =& -\mathrm{i}\omega \mathbf{D}(\mathbf{r},\omega) + \mathbf{j}(\mathbf{r},\omega), & (1.12)\\ \nabla \cdot \mathbf{D}(\mathbf{r},\omega) =& \rho(\mathbf{r},\omega), & (1.13)\\ \nabla \cdot \mathbf{B}(\mathbf{r},\omega) =& 0, & (1.14)\end{aligned}$$

---

[1] In the remainder of this thesis, I will often omit the attribute "angular" and refer to ω as a frequency, although this is slightly inaccurate.

with the corresponding constitutive relations:

$$\mathbf{D}(\mathbf{r},\omega) = \epsilon_0\epsilon(\mathbf{r})\mathbf{E}(\mathbf{r},\omega), \tag{1.15}$$

$$\mathbf{B}(\mathbf{r},\omega) = \mu_0\mathbf{H}(\mathbf{r},\omega). \tag{1.16}$$

Maxwell's curl equations (Eqn. (1.1) and Eqn. (1.2)) are two partial differential equations of first order that connect four different fields. By inserting one of the curl equation into the other one and with the help of the constitutive relations, we can derive a single wave equation for one type of field.

Throughout this thesis, I will focus on the magnetic wave equation. To derive this equation, we first replace the electric displacement field **D** in Eqn. (1.2) according to Eqn. (1.5):

$$0 = \nabla\times\mathbf{H}(\mathbf{r},t) - \partial_t\mathbf{D}(\mathbf{r},t) - \mathbf{j}(\mathbf{r},t) \tag{1.17}$$

$$= \epsilon^{-1}(\mathbf{r})\nabla\times\mathbf{H}(\mathbf{r},t) - \epsilon_0\partial_t\mathbf{E}(\mathbf{r},t) - \epsilon^{-1}(\mathbf{r})\mathbf{j}(\mathbf{r},t). \tag{1.18}$$

As this is identically zero for every **r**, the same holds for the curl[2]:

$$\nabla\times\Big(\epsilon^{-1}(\mathbf{r})\nabla\times\mathbf{H}(\mathbf{r},t) - \epsilon_0\partial_t\mathbf{E}(\mathbf{r},t) - \epsilon^{-1}(\mathbf{r})\mathbf{j}(\mathbf{r},t)\Big) = 0. \tag{1.19}$$

We obtain the driven wave equation by eliminating the parentheses and inserting Eqn. (1.1) and Eqn. (1.6):

$$\nabla\times\Big(\epsilon^{-1}(\mathbf{r})\nabla\times\mathbf{H}(\mathbf{r},t)\Big) - \epsilon_0\partial_t\nabla\times\mathbf{E}(\mathbf{r},t) - \nabla\times\Big(\epsilon^{-1}(\mathbf{r})\mathbf{j}(\mathbf{r},t)\Big) = 0, \tag{1.20}$$

$$\nabla\times\Big(\epsilon^{-1}(\mathbf{r})\nabla\times\mathbf{H}(\mathbf{r},t)\Big) + \epsilon_0\mu_0\partial_t^2\mathbf{H}(\mathbf{r},t) = \nabla\times\Big(\epsilon^{-1}(\mathbf{r})\mathbf{j}(\mathbf{r},t)\Big), \tag{1.21}$$

where the right hand side in Eqn. (1.21) is derived from the current distribution of the source. Naively, one might wonder[3] if this term is sufficient for all possible sources or if an additional term containing the charge distribution might be required under certain circumstances. This is not necessary if the electric charge density ρ and the current **j** are connected via the continuity equation

$$\nabla\mathbf{j} + \partial_t\rho = 0. \tag{1.22}$$

Eqn. (1.21) is the master equation for the magnetic component of an electrodynamic problem in the sense that every physical magnetic field is a solution. Accordingly, Eqn. (1.21) describes the time evolution of any physical field. However, not every solution of Eqn. (1.21) represents a physical field. The solution space is restricted by the absence

---

[2] At this point we replace an equation with an expression for its spatial derivative. As the inverse operation is not unique, this step might increase the solution space for the equation. Indeed, the curl derivative allows for adding any irrotational field to the original equation. This is fixed by Eqn. (1.3) and Eqn. (1.22).

[3] At least I did at some point.

of magnetic monopoles, which is formulated in Eqn. (1.4). This equation is a side condition for the physical solutions of the wave equation and only allows for solutions with divergence–free magnetic field.

If we omit the source term of Eqn. (1.21) and apply a time–harmonic ansatz (Eqn. (1.7)) to the magnetic field, we obtain the undriven stationary wave equation

$$\nabla \times \Big(\epsilon^{-1}(\mathbf{r})\nabla \times \mathbf{H}(\mathbf{r})\Big) - \frac{\omega^2}{c^2}\mathbf{H}(\mathbf{r}) = 0, \tag{1.23}$$

with $c = (\epsilon_0\mu_0)^{-1/2}$ being the speed of light in vacuum. This is an eigenvalue equation and together with the magnetic divergence condition Eqn. (1.4) it provides the (stationary) eigenmodes of a system.

A master equation similar to Eqn. (1.21) formulated in terms of the electric field can be derived by inserting Eqn. (1.2) into the curl of Eqn. (1.1):

$$\epsilon^{-1}(\mathbf{r})\nabla \times \nabla \times \mathbf{E}(\mathbf{r}) + \frac{\partial_t^2}{c^2}\mathbf{E}(\mathbf{r}) = -\epsilon^{-1}(\mathbf{r})\partial_t\mathbf{j}, \tag{1.24}$$

The corresponding stationary eigenvalue problem is:

$$\epsilon^{-1}(\mathbf{r})\nabla \times \nabla \times \mathbf{E}(\mathbf{r}) - \frac{\omega^2}{c^2}\mathbf{E}(\mathbf{r}) = 0. \tag{1.25}$$

This formulation has two disadvantages as compared to Eqn. (1.21). First, the corresponding divergence condition

$$\nabla \cdot \Big(\epsilon(\mathbf{r})\mathbf{E}(\mathbf{r})\Big) = \frac{\rho(\mathbf{r})}{\epsilon_0} \tag{1.26}$$

is considerably more complicated, notably in the presence of a current density. Second, the magnetic wave operator $\nabla \times (\epsilon^{-1}(\mathbf{r})\nabla\times)$ is Hermitian with respect to the conventional L^2 scalar product whereas the electric operator $\epsilon^{-1}(\mathbf{r})\nabla \times \nabla\times$ is not[4]. However, Hermiticity is a numerically very desirable property as it usually allows to solve the problem at hand with specialized algorithms that require less memory and computational time.

Finally, in numerical simulations, it is very convenient to express all quantities in a set of physical units such that the speed of light becomes one. This can be accomplished by rescaling the magnetic field and time with appropriate factors

$$\begin{aligned} \mathbf{H}'(\mathbf{r},t') &= \sqrt{\frac{\epsilon_0}{\mu_0}}\mathbf{H}(\mathbf{r},t) && (1.27)\\ t' &= \sqrt{\frac{1}{\epsilon_0\mu_0}}t, && (1.28) \end{aligned}$$

---

[4] Hermiticity implies real–valued eigenvalues and an orthonormal eigensystem. The spectra of the magnetic and the electric wave equation are identical but the field distributions of the corresponding eigenmodes are not. The magnetic fields fulfill the conventional orthonormality condition $\int \mathrm{d}r\mathbf{H}^*_\alpha(\mathbf{r})\mathbf{H}_\beta(\mathbf{r}) = \delta_{\alpha\beta}$, whereas the “orthogonality” condition for the electric fields contains an additional integral kernel $\int \mathrm{d}r\mathbf{E}^*_\alpha(\mathbf{r})\epsilon(\mathbf{r})\mathbf{E}_\beta(\mathbf{r}) = \delta_{\alpha\beta}$. Here, α and β denote two eigenmodes.

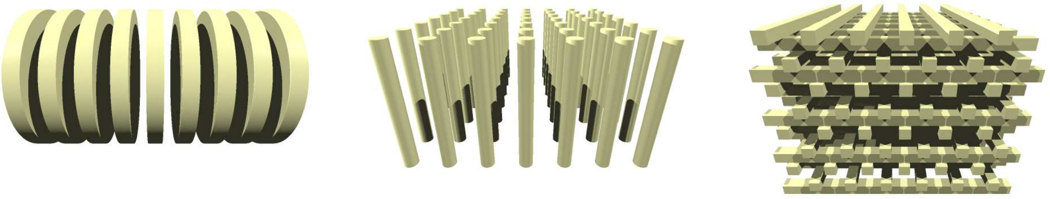
(a) Bragg stack (b) Post array (c) Woodpile

Figure 1.1.: *Examples for one–dimensional (a), two–dimensional (b) and three–dimensional (c) photonic crystals.*

while the electric field is not modified:

$$\mathbf{E}'(\mathbf{r}, t') = \mathbf{E}(\mathbf{r}, t). \tag{1.29}$$

In this transformed unit system, time has the same unit as distance and, as desired, the speed of light vanishes from all equations:

$$c' = 1. \tag{1.30}$$

1.2. Photonic Crystals

Photonic crystals (PhC) are dielectric structures which exhibit a periodicity on the length scale of the wave length of light. The periodicity can be present in one, in two and in three independent directions (cf. Fig. 1.1).

1.2.1. Periodicity

A photonic crystal is a periodic structure which can be constructed by replicating a small structure (basis) on each site of a Bravais lattice. This lattice is described by a set of one, two or three primitive lattice vectors for a one–, two– or three–dimensional lattice, respectively. Conventionally, they are denoted $\mathbf{a}_i$. Each of the lattice sites $\mathbf{R}$ is the sum of integer multiples of the lattice primitives:

$$\mathbf{R} = \sum_{i=1,2,3} r_i \mathbf{a}_i, \quad r_i \in \mathbb{Z}. \tag{1.31}$$

These vectors are also known as lattice vectors. The area which is closer to one specific lattice site than to any other is called a Wigner–Seitz cell (WSC; cf. Fig. 1.2 for an illustration in two dimensions).

The material properties of the crystal and thus the wave equation are periodic with respect to the lattice, i. e. invariant under translation by a lattice vector:

$$\epsilon(\mathbf{r}) = \epsilon(\mathbf{r} + \mathbf{R}). \tag{1.32}$$

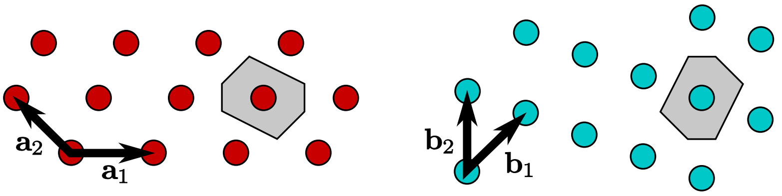

Figure 1.2.: *Sketch of a 2D Bravais lattice (lattice sites indicated by red circles) and the corresponding reciprocal lattice (indicated by blue circles). The lattice primitives and reciprocal primitives are denoted by* $\mathbf{a}_{1/2}$ *and* $\mathbf{b}_{1/2}$*, respectively. The Wigner–Seitz cell and first Brillouin zone are highlighted in gray.*

As in the case of an electronic crystal one consequence of this periodicity is the fact that the eigenstates fulfill Bloch's theorem [20, 21]. They are invariant under lattice translation except for a plane wave phase factor:

$$\mathbf{H}(\mathbf{r}) = \exp\left(\mathrm{i}\mathbf{k}\mathbf{r}\right)\mathbf{h}_{\mathbf{k}}(\mathbf{r}) \quad \text{with} \quad \mathbf{h}_{\mathbf{k}}(\mathbf{r}) = \mathbf{h}_{\mathbf{k}}(\mathbf{r}+\mathbf{R}), \tag{1.33}$$

where the lattice periodic parts $\mathbf{h}_{\mathbf{k}}(\mathbf{r})$ for one wave vector $\mathbf{k}$ are solutions to the equation

$$(\nabla + \mathbf{k}) \times \left(\epsilon^{-1}(\mathbf{r})(\nabla + \mathbf{k}) \times \mathbf{h}(\mathbf{r})\right) - \frac{\omega^2}{c^2}\mathbf{h}(\mathbf{r}) = 0. \tag{1.34}$$

The space of wave vectors $\mathbf{k}$, the reciprocal space, is also invariant under the translations of a Bravais lattice (the reciprocal lattice) that is spanned by the reciprocal primitives $\mathbf{b}_i$ defined by

$$\mathbf{b}_i \cdot \mathbf{a}_j = 2\pi\delta_{ij}, \tag{1.35}$$

where δ_{ij} is the Kronecker symbol. Wave vectors that are connected by a reciprocal lattice vector

$$\mathbf{G} = \sum_{i=1,2,3} g_i \mathbf{b}_i, \quad g_i \in \mathbb{Z}, \tag{1.36}$$

are equivalent in the sense that they correspond to the same eigenvalues and eigenfunctions of the wave equation. The Wigner–Seitz cell of this reciprocal lattice is usually referred to as the first Brillouin zone (BZ) (cf. Fig. 1.2).

Finally, the physical system can be subject to point symmetries such as rotations or inversion with respect to $\mathbf{r} = \mathbf{0}$ or a combination of point operations and fractional lattice translations. All such operations that leave the crystal invariant form a group with respect to concatenation: The space group $\mathcal{S}$. The point–symmetry parts of all space group elements form another group themselves [20], the point group $\mathcal{P}$. However, the members of $\mathcal{P}$ not necessarily leave the system invariant[5]. Each wave vector $\mathbf{k}$ in the first Brillouin zone belongs to one set $\{\mathbf{k}\}$ that is invariant under all elements of $\mathcal{P}$ and

---

[5] They do if $\mathcal{S}$ is symmorphic, i. e. , if $\mathcal{S}$ can be constructed from operations that only contain integer lattice translations. In this (and only in this) case is $\mathcal{P} \subset \mathcal{S}$.

referred to as the star of $\mathbf{k}$. The sets of eigenvalues for wave vectors on the same star are identical and the eigenmodes are connected via the respective symmetry operation and a unitary basis transformation (via the unitary representation matrix [20, 22]).

If the PhC features a two–dimensional lattice and is homogeneous in the perpendicular direction (conventionally denoted z), it is usually referred to as a two–dimensional photonic crystal. If furthermore the propagation of light is restricted to the plane of periodicity, the solutions can be divided into two classes according to their field orientation (polarization). For one half of the physical modes, the magnetic field component is solely oriented along the z–axis (H–polarization) and the modes are invariant under any translation along this direction. The other physical modes feature an electric field with only the z–component (E–polarization). Both subspaces are decoupled and can are described individually by the scalar wave equations

$$\epsilon^{-1}(\mathbf{r})(\partial_x^2 + \partial_y^2)E_z(\mathbf{r},t) - \frac{\partial_t^2}{c^2}E_z(\mathbf{r},t) = 0 \qquad \text{(E–polarization)} \tag{1.37}$$

$$\begin{pmatrix} \partial_x \\ \partial_y \end{pmatrix} \cdot \left[\epsilon^{-1}(\mathbf{r}) \begin{pmatrix} \partial_x \\ \partial_y \end{pmatrix} H_z(\mathbf{r},t) \right] - \frac{\partial_t^2}{c^2}H_y(\mathbf{r},t) = 0 \qquad \text{(H–polarization)} \tag{1.38}$$

It is noteworthy that these equations do not feature any unphysical solutions (the divergence conditions are always implicitly fulfilled). Hence, the respective solution spaces form complete bases.

1.2.2. Band Structure

The periodicity of reciprocal space implies that the dispersion relation of a photonic crystal is periodic with respect to $\mathbf{k}$. In analogy to the electric potential in an electronic crystal [20, 23], the non–vanishing periodic modulation of the dielectric distribution results in a dispersion relation that is deformed with respect to that of a homogeneous medium. In particular, certain degeneracies are split, e. g. , at the border of the Brillouin zone and at the Γ–point. This leads to the formation of frequency ranges in which the propagation of light in a particular direction is prohibited, the so–called stop bands. For certain geometries such as the woodpile structure [24], all stop bands can be brought to overlap in a common frequency range for a sufficiently high refractive index contrast (the ratio between the refractive indices of the PhCs constituent materials) Consequently, propagation of light with a frequency inside this range is prohibited in any direction that does not leave the periodicity[6]. Such a frequency range is called a complete photonic band gap (PBG) and features a complete lack of photonic states, which results in interesting physical effects such as a strong suppression of spontaneous emission (cf. Ref. [1] but also Chap. 5).

Another implication of the deformed dispersion relation is the fact that group and phase velocity become independent of each other to a very high degree which results

---

[6] In a Bragg stack, this is only the stacking direction; for a 2D photonic crystal, this covers all directions in the plane of periodicity.

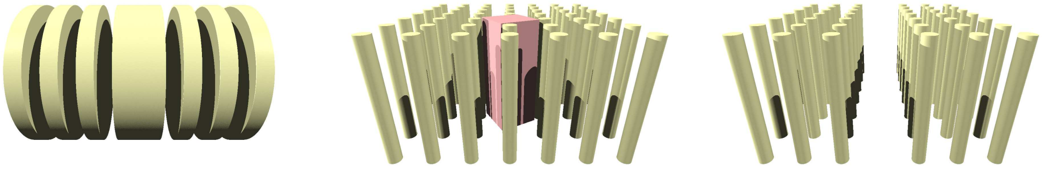

(a) Cavity in Bragg stack (b) Cavity in post array (c) Wave guide in post array

Figure 1.3.: *Sketch of zero–dimensional (a) and (b) defects and of a one–dimensional one (c). We refer to point–like defects with localized states as cavities even if the mode volume is filled with material as in (a) and (b).*

in highly anisotropic energy propagation inside the crystal [25]. Most notably is the occurrence of regions with vanishing group velocity. Here, light seems to acquire an effective mass and the interaction between light and matter is increased due to a pile–up effect. This slow light effect always occurs at the edges of any stop band (most notably at the PBG edges). However, it is not restricted to such spectral regions.

Finally, the plane waves known to be the propagating waves in homogeneous media cease to be solutions of the wave equation in a photonic crystal. Instead, light travels as Bloch modes which may closely resemble plane waves, e. g. for very small frequencies (the so–called effective material limit) or very weak index contrast. For sufficiently high frequencies and greater index contrast, however, well–known properties of plane waves such as polarization[7] may no longer be well–defined. We will come back to this in Chap. 4.

1.2.3. Breaking Periodicity

Periodicity is the basic prerequisite for the formation of a spectral range without photonic states (i. e. a complete band gap). Thus, a properly designed modification of the crystal introduces states inside this spectral range. If the modification is local or extended, we may expect the corresponding modes to be similarly localized or extended, respectively. Some examples for such defects are depicted in Fig. 1.3.

We may classify defects by how many independent discrete translational symmetries of the host crystal's lattice remain. If the system including the defect does not retain any periodicity and if the defect is localized, we refer to this as a zero–dimensional defect or a cavity[8]. Inside the band gap of the host crystal, such cavities exhibit a discrete spectrum and exponentially localized eigenmodes. Photonic crystal cavities are known

---

[7] Only in the case of a 3D photonic crystal

[8] The term "cavity" suggests a hollow space. However, we also use this term for defects that consists of high–index material.

for exhibiting extraordinarily high quality factors [18] and low mode volume which makes them a very popular system for investigating the interaction of light and single electronic systems [26] and low threshold lasing.

If the the defect retains one periodic direction of the host crystal, the whole structure including the defect forms a one–dimensional PhC with its own dispersion relation. The reciprocal space is restricted to the interval $[-\frac{\pi}{a}, \frac{\pi}{a}]$ on a line parallel to the direction of periodicity, where a is the lattice constant of the wave guide. According to the Bloch theorem, the eigenmodes of such a system are modulated plane waves. If the frequency lies inside the host crystal's band gap[9]. the modulation decays exponentially away from the line defect. If such a line defect supports confined propagating modes along one direction, it is usually referred to as a photonic crystal wave guide. PhC wave guides provide the possibility to guide light on a length scale much smaller than possible with index guiding as employed, e. g. , in optical fibers. As they are a special case of a photonic crystal, they may also provide features such as a regime with vanishing group velocity.

A two–dimensional defect, finally, forms a two–dimensional photonic crystal that can, for instance, serve as a host crystal for cavities or wave guides. Examples for this are thin two–dimensional PhC structures wrapped in a 3D PhC with a complete band gap [27] or (more mundane and without confined modes) stacking faults that occur in the self–organized growth process of opal films.

1.2.4. Wave Guide Circuits

The zero– and one–dimensional defects introduced in the previous section can be used as resonators and wave guides in a photonic crystal, respectively. These basic functional elements together with wave guide bends (cf. Fig. 1.4b) can be combined to form photonic circuits with complex functionality. Fig. 1.4 shows possible implementations of a narrow–band filter based on a point cavity and of a directional coupler formed by wave guides and wave guide bends. Finally, these building blocks of moderate complexity can be combined to form large scale circuits such as an electrically tunable Mach–Zehnder interferometer [31].

During their PhDs, Matthias Schillinger [29] and Daniel Hermann [30] developed methods to simulate such functional devices and complex photonic circuits based on a expression in terms of photonic Wannier functions and studied possibilities to systematically optimize them. These successful works are the motivation for me further developing the Wannier approach (cf. Sec. 3.5) and my attempt to extend it to three–dimensional systems (cf. Sec. 3.4).

---

[9] More precisely, if the host PhCs dispersion relation projected onto the wave guide's Brillouin zone does not provide any propagating states for the wave vector and frequency of the line defect mode. This is necessarily satisfied for frequencies inside the band gap.

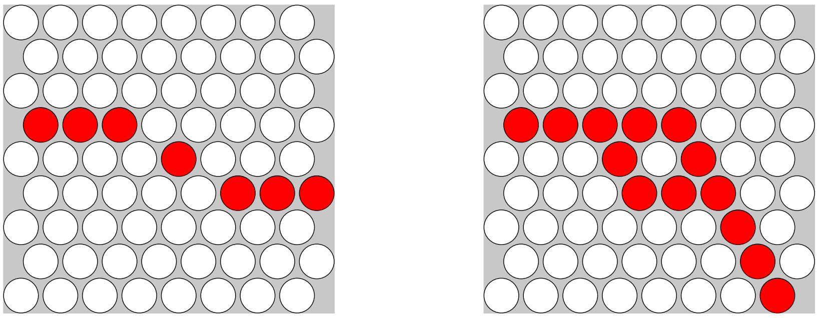

(a) narrow band–pass filter

(b) broad–band wave guide bend

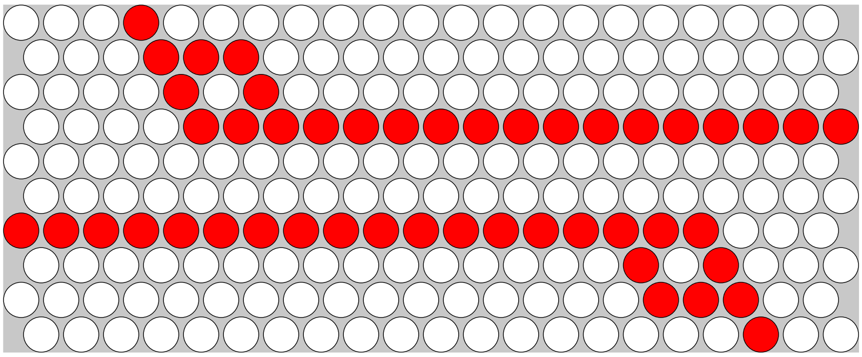

(c) wave guide directional coupler

Figure 1.4.: *Schematic examples for functional building blocks inscribed into a 2D host crystal consisting of air pores (relative radius* $r/a = 0.45$*) in silicon. The red circles indicate defect sites where the air hole is filled with a dielectric material, e. g. , by infiltration of individual pores [28]. The narrow band–pass filter (a) actually consists of three basic elements: Two wave guides and a single cavity resonator in between. The wave guide bend (b) allows to modify the direction of propagation of guided modes. Finally, two wave guide bends and four wave guides can be arranged to form a coupler. These sketches were motivated by optimized designs presented in Refs. [29, 30].*

2

Chapter 2.

Band Structure

The photonic band structure plays a central role for characterizing the properties of a photonic crystal (PhC). For investigating real–world systems such as woodpile structures, it is inevitable to compute the band structure numerically.

In Sec. 2.1, I give an introduction to the conventional band structure problem and briefly address two popular ways to solve it numerically. Then I present some basic thoughts about the occurrence of stationary points in a band structure (Sec. 2.2). Finally, the chapter ends with a brief introduction to the class of on–shell band structure calculations (Sec. 2.3) and I qualitatively describe two possible ways to solve this type of problems numerically.

2.1. Conventional Band Structure — Task and Solution Properties

The band structure of a photonic crystal is the dispersion relation of all propagating modes (i.e. real wave vector and real frequency)[1]. As it relates wave vectors inside the first Brillouin zone to allowed frequencies, the numerical methods for its calculation can be roughly divided into two categories. The first possibility is to regard the wave vector as a parameter and calculate the allowed frequencies for a set of $\mathbf{k}$–vectors. This "conventional" band structure approach is the topic of this section. The second possibility is to consider the frequency as a parameter and solve for allowed wave vectors. Sec. 2.3 deals with this so–called on–shell approach.

The dispersion relation of continuous waves in a periodic medium typically consists of infinitely many branches. At each $\mathbf{k}$–point, the allowed frequencies ω are enumerated in ascending order. All pairs $(\mathbf{k};\omega)$ that are assigned the same number form one *band* of the dispersion relation. The corresponding band number is usually denoted by n. A branch that is not connected to any other branch always forms an individual band.

---

[1] This definition can be extended to include the evanescent modes as well

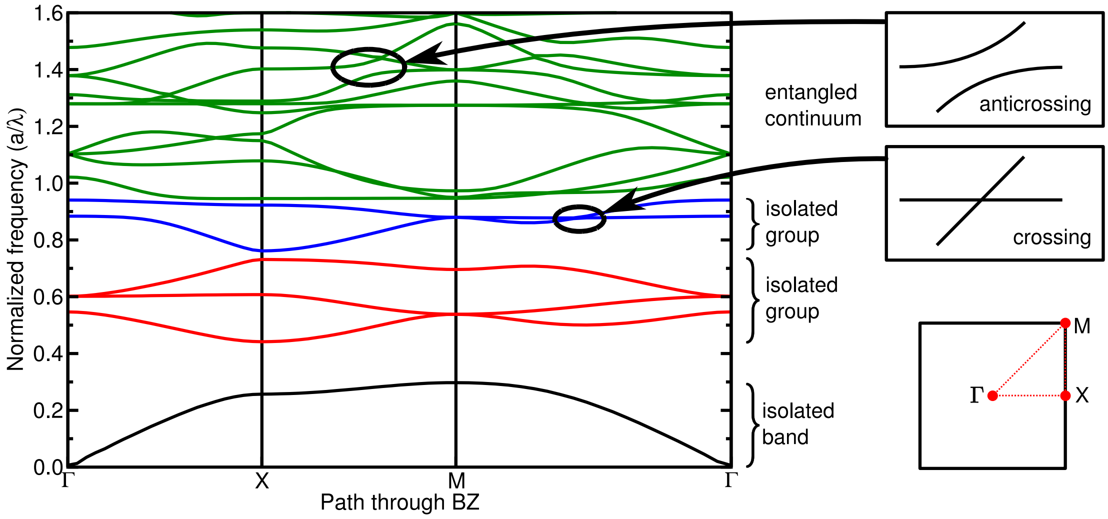

Figure 2.1.: *Example to illustrate several features of a typical band structure plot. The system for which this plot has been calculated is a square lattice of cylindrical silicon posts (radius* $r = 0.18$ *times the lattice constant* a*,* $\epsilon = 12$*) in E–polarization. It is common to plot the band structure of two– and three–dimensional crystals on a specific path of the first Brillouin zone (right bottom panel) between certain high–symmetry points as all interesting features are expected to be visible there. This plot features three isolated band groups (which in turn consists of entangled bands except for the first band, which is isolated) and a continuum of entangled bands. For illustrative purposes, I highlighted a band crossing and a particularly beautiful anticrossing. The calculation was performed using the plane wave expansion.*

Each band covers a certain frequency range. If these frequency ranges overlap or if the bands have pairs $(\mathbf{k};\omega)$ in common (are degenerate at some $\mathbf{k}$–point), these bands are said to be *entangled*. A single band or several bands forming one group that are not entangled with other bands are called an isolated band or an isolated group, respectively. A point of degeneracy is called a band crossing. In contrast, two bands that approach each other as if they will cross but avoid the crossing form a so–called anticrossing. All these features are illustrated in Fig. 2.1 at the example of a 2D photonic crystal.

2.1.1. Plane Wave Expansion

The most straightforward way to solve a linear wave equation with a periodic potential is by switching to momentum space. The great advantage is that spatial derivatives are effectively translated to diagonal matrices. The periodic potential translates to a (typically dense) matrix that consists of the Fourier components of the potential.

Finally, diagonalizing the system matrix gives the eigenfrequencies and eigenmodes of the system. In solid state theory, this plane wave expansion (PWE) has been a standard approach for decades and it is well–established in photonics as well. In the following, I will provide a very brief review of the plane wave method for photonic crystals based on Ref. [32]. We start with the wave equation in terms of the magnetic field distribution:

$$\nabla \times \Big[\eta(\mathbf{r})\Big(\nabla \times \mathbf{H}(\mathbf{r})\Big)\Big] = \frac{1}{c^2}\partial_t^2 \mathbf{H}(\mathbf{r}), \tag{2.1}$$

where $\eta(\mathbf{r}) = \epsilon^{-1}(\mathbf{r})$ is the inverse dielectric function. An alternative formulation in terms of the electric field is possible and particularly useful for certain two–dimensional systems. However, throughout my thesis, I only use the magnetic field distributions of Bloch modes and, thus, forgo a presentation of the electric wave equation.

The stationary solutions (i. e. those associated to a specific frequency) of this equation are Bloch modes. According to the Bloch–Floquet theorem [20], we may decompose this into the product of a plane wave and a periodic function $\mathbf{h}_{n\mathbf{k}}(\mathbf{r})$, which is periodic with respect to $\mathbf{r}$ and $\mathbf{k}$:

$$\mathbf{H}_{n\mathbf{k}}(\mathbf{r}) = \mathbf{h}_{n\mathbf{k}}(\mathbf{r}) \exp\big(-\mathrm{i}\mathbf{k}\mathbf{r}\big)\,. \tag{2.2}$$

Here, $\mathbf{k}$ is the wave vector and n is the band index. Typically, n sorts the solutions with respect to frequency in ascending order. As $\mathbf{h}_{n\mathbf{k}}(\mathbf{r})$ is periodic with respect to the crystal's lattice, we may expand it into a Fourier series:

$$\mathbf{h}_{n\mathbf{k}}(\mathbf{r}) = \sum_{\mathbf{G}} \mathbf{h}_{n\mathbf{k}}^{(\mathbf{G})} \exp\big(-\mathrm{i}\mathbf{G}\mathbf{r}\big)\,, \tag{2.3}$$

$$\mathbf{H}_{n\mathbf{k}}(\mathbf{r}) = \sum_{\mathbf{G}} \mathbf{h}_{n\mathbf{k}}^{(\mathbf{G})} \exp\big(-\mathrm{i}(\mathbf{k}+\mathbf{G})\mathbf{r}\big)\,, \tag{2.4}$$

where $\mathbf{h}_{n\mathbf{k}}^{(\mathbf{G})}$ are the vector–valued Fourier coefficients and the summation runs over all reciprocal lattice vectors. Using this discretization and the fact that we aim for solutions with well–defined frequencies, we can explicitly evaluate the derivatives in Eqn. (2.1) using the simple translation rules:

$$\partial_t \rightarrow \mathrm{i}\omega, \tag{2.5}$$

$$\partial_{r_j} \rightarrow -\mathrm{i}(k_j + G_j). \tag{2.6}$$

As we effectively perform a Fourier transform, the product between η and the curl of $\mathbf{H}$ is transformed to a convolution:

$$\sum_{\mathbf{G}'} \eta_{\mathbf{G}-\mathbf{G}'}(\mathbf{k}+\mathbf{G}) \times (\mathbf{k}+\mathbf{G}') \times \mathbf{h}_{n\mathbf{k}}^{(\mathbf{G}')} = \frac{\omega^2}{c^2}\mathbf{h}_{n\mathbf{k}}^{\mathbf{G}}, \tag{2.7}$$

where the Fourier coefficients of the inverse dielectric distribution are given by[2]

$$\eta_{\mathbf{G}-\mathbf{G}'} = \int_{\mathrm{WSC}} \mathrm{d}r \eta(\mathbf{r}) \exp\Big(\mathrm{i}(\mathbf{G}-\mathbf{G}')\mathbf{r}\Big)\,. \tag{2.10}$$

---

[2] The convergence is improved if this naive representation of the inverse dielectric function is replaced

Each Fourier coefficient $\mathbf{h}_{n\mathbf{k}}^{(\mathbf{G})}$ of the field contains three degrees of freedom: the three vector components. We can use the divergence condition for the magnetic field to reduce this number by one:

$$\begin{aligned}
0 &= \nabla \cdot \left[\mathbf{h}_{n\mathbf{k}}^{(\mathbf{G})} \exp\left(-\mathrm{i}(\mathbf{k}+\mathbf{G})\mathbf{r}\right) \right] && (2.11)\\
&= (\mathbf{k}+\mathbf{G}) \cdot \mathbf{h}_{n\mathbf{k}}^{(\mathbf{G})} \exp\left(-\mathrm{i}(\mathbf{k}+\mathbf{G})\mathbf{r}\right), && (2.12)\\
\Rightarrow \mathbf{h}_{n\mathbf{k}}^{(\mathbf{G})} &= h_{n\mathbf{k}}^{(\mathbf{G},1)} \hat{e}^{(\mathbf{G},1)} + h_{n\mathbf{k}}^{(\mathbf{G},2)} \hat{e}^{(\mathbf{G},2)}, && (2.13)
\end{aligned}$$

where $\hat{e}^{(\mathbf{G},1)}$ and $\hat{e}^{(\mathbf{G},1)}$ are unit vectors orthogonal to each other and to $(\mathbf{k}+\mathbf{G})$. The final eigenvalue problem is:

$$\begin{aligned}
\sum_{\mathbf{G}'} |\mathbf{k}+\mathbf{G}||\mathbf{k}+\mathbf{G}'|\eta_{\mathbf{G}-\mathbf{G}'} \begin{pmatrix} \hat{e}^{(\mathbf{G}),2}\hat{e}^{(\mathbf{G}'),2} & -\hat{e}^{(\mathbf{G}),2}\hat{e}^{(\mathbf{G}'),1} \\ -\hat{e}^{(\mathbf{G}),1}\hat{e}^{(\mathbf{G}'),2} & \hat{e}^{(\mathbf{G}),1}\hat{e}^{(\mathbf{G}'),1} \end{pmatrix} \begin{pmatrix} h_{n\mathbf{k}}^{(\mathbf{G}'),1} \\ h_{n\mathbf{k}}^{(\mathbf{G}'),2} \end{pmatrix} & \\
= \frac{\omega^2}{c^2} \begin{pmatrix} h_{n\mathbf{k}}^{(\mathbf{G}),1} \\ h_{n\mathbf{k}}^{(\mathbf{G}),2} \end{pmatrix}. & \quad (2.14)
\end{aligned}$$

The summation still runs over all infinitely many reciprocal lattice vectors. For this to be solvable on a finite computer, the Fourier expansions of the field and the dielectric function must be restricted to a common finite subset of the reciprocal lattice. Thereby, we obtain a matrix problem, which can be numerically diagonalized, e. g. , using the numeric library `LAPACK`. Usually, only the lowest bands are of interest which suggests to calculate only those eigenvalues by means of an iterative solver such as the Arnoldi algorithm.

Excluding the unphysical solutions of the wave equation (cf. Eqn. (2.13)) has mainly positive effects. First, it reduces the matrix size from $3N \times 3N$ to $2N \times 2N$ for a given number N of reciprocal lattice vectors $\mathbf{G}$. Second, all unphysical solutions have eigenfrequency zero and thus form an in principle infinitely large degenerate subspace that would cause difficulties if it was not explicitly excluded: For instance, the vast number of modes with frequency zero would make it rather difficult to use iterative eigensolvers to compute the lower bands[3]. Furthermore, these unphysical modes would cause major trouble in the generation of Wannier functions (cf. Chap. 3).

As a photonic crystal typically consists of regions with constant permittivities and discontinuities at any border between two such regions, the Fourier transform of η typically decays like $|\mathbf{G}|^{-1}$. As a consequence, the matrix problem is dense and exhibits only first order convergence with respect to the maximum value of $|\mathbf{G}|$.

---

with the inverse of the convolution–matrix of the dielectric distribution:

$$\begin{aligned}
\eta &= \epsilon^{-1}, && (2.8)\\
\epsilon_{\mathbf{G},\mathbf{G}'} &= \int_{\mathrm{WSC}} \mathrm{d}r\epsilon(\mathbf{r}) \exp\left(\mathrm{i}(\mathbf{G}-\mathbf{G}')\mathbf{r}\right). && (2.9)
\end{aligned}$$

[3] Calculation e. g. of defect modes in a band gap would not be affected, though.

2.1.2. The MPB Package

The plane wave expansion suffers from some performance problems. First, the convergence rate is limited to first order due to the discontinuous dielectric function. Second, the matrix to be diagonalized is dense. This implies that the matrix–vector product scales quadratically and a full diagonalization cubically in time. Third, if $\eta_{\mathbf{G},\mathbf{G}'}$ is calculated via a matrix inversion (cf. Eqn. (2.9)), Krylov–based eigensolvers such as the Arnoldi–method do not improve the scaling of the computational cost as the computation will be dominated by the matrix inversion.

These three limitations are overcome by a method developed and implemented by S. Johnson [16]. This approach is conceptually similar to the plane wave expansion as it eliminates the spatial derivatives by a Fourier transform of the electromagnetic field. Furthermore, the number of degrees of freedom is reduced from $3N$ to $2N$ by excluding longitudinally polarized basis functions (cf. Eqn. (2.13)). Here N is the number of Fourier coefficients[4] taken into account.

In contrast to the PWE, however, the convolution in Eqn. (2.7) is replaced with a Fourier transform, a simple real–space multiplication and an inverse Fourier transform. This is favorable because the fast Fourier transform algorithm only scales with $\mathcal{O}(n \log n)$. Furthermore, it is very memory–efficient because no system matrix has to be stored explicitly. This also suggests to use a Krylov–type eigensolver, which just requires the action of the operator on a given state rather than the explicit matrix representation. Finally, the multiplication with the dielectric function in real–space allows to improve the accuracy of the spatial discretization to second order.

Putting all together, the approach sketched above features superior performance as compared to the standard plane wave expansion at good accuracy. The reference implementation [16] by S. Johnson is known as the MIT Photonic Bands (MPB) package and very popular in photonic crystal research. Most band structure calculations which I performed for this thesis were obtained with this program.

2.2. Stationary Points of Photonic Bands

The numerical methods outlined in Sec. 2.1 provide the band structure, i. e. , the eigenmodes and eigenfrequencies for each individual point in reciprocal space. In practice, band structure calculations are often restricted to the outline of the irreducible Brillouin zone (iBZ). This is a minimal set of $\mathbf{k}$–points that reproduces the full first Brillouin zone if transformed according to every element of the crystal's point group. Usually, it is chosen to be a wedge–like shape bounded by high–symmetry lines. Although there is no formal justification, band extrema are generally assumed to occur on these lines of high symmetry. Curiously, this holds at least approximately true in many practical situations, although it has been shown that in general, band extrema can occur anywhere in the first Brillouin zone [33].

---

[4] The Fourier expansion within the MPB method should not be confused with a true plane wave expansion as the dielectric function is not translated to Fourier space.

However, the systems investigated in Ref. [33] are highly asymmetric discrete graphs or continuous structures where the unit cell contains large definition gaps, i. e. , is not simply connected. Such systems are not extraordinarily similar to typical photonic crystals.

In this section, I would like to provide some ideas that might help judging if a given structure has its maxima on the high symmetry lines or not. First, I discuss which symmetries of the crystal for sure lead to extrema at high symmetry points without excluding extrema somewhere else. In the second part, I employ the Wannier basis to estimate an upper bound for the total number of extrema of a band or a band group. For simplicity of illustration I mostly restrict myself to 2D systems based on the square and triangular lattices. Furthermore, it is advantageous to consider eigenvalue bands

$$\lambda_{n\mathbf{k}} = \frac{\omega_{n\mathbf{k}}^2}{c^2} \tag{2.15}$$

rather than the frequency bands $\omega_{n\mathbf{k}}$ themselves.

The relevant wave operators are quadratic (thus holomorphic) in $\mathbf{k}$. Hence, the eigenvalues away from any $\mathbf{k}$–points of degeneracy[5] form holomorphic functions of $\mathbf{k}$ [33] (see also Sec. 2.2.3 for a finite number of bands). The consequences of this holomorphy include: An eigenvalue band that is not degenerate at some point P is also not degenerate in a neighborhood of this point. I would like to call such a band *solitary* at P. A solitary eigenvalue band is smooth in the whole neighborhood. If the eigenvalue band is smooth and positive in the whole neighborhood, its square root (i.e. the frequency band) is smooth, too. In the remainder, I will refer to eigenvalue bands instead of frequency bands, because this avoids handling the first bands at Γ separately.

Four facts are crucial [20]. First, the eigenvalues associated with $\mathbf{k}$–points that are connected by symmetry operations of the crystal's point group (the star of $\mathbf{k}$) are identical. Second, the eigenvalues associated with $\mathbf{k}$–points that are connected by a reciprocal lattice translation (equivalent $\mathbf{k}$–points) are identical, irrespective of point symmetry. Third, I generally assume time–reversal symmetry and, as a consequence, inversion symmetry in reciprocal space of the eigenvalue bands. Finally, I consider solitary bands[6], only. This is justified by the observation, that in practice degenerate extrema rarely define a band gap[7].

2.2.1. Existence of Band Extrema at High Symmetry Points

First, I discuss stationary points induced by periodicity and time–reversal alone. Inversion symmetry in reciprocal space connects $\mathbf{k}$ and $-\mathbf{k}$; consequently, their eigenvalues

---

[5] In one dimension it may be possible to define smooth generalized bands by appropriately choosing the band labels left and right of a point of degeneracy. However, in two (and consequently in three) dimensions this is not possible for discrete points of degeneracy because "left" and "right" of the point of degeneracy cannot be defined.

[6] This includes bands that are degenerate somewhere except for the point of consideration.

[7] At least I am not aware of any such case.

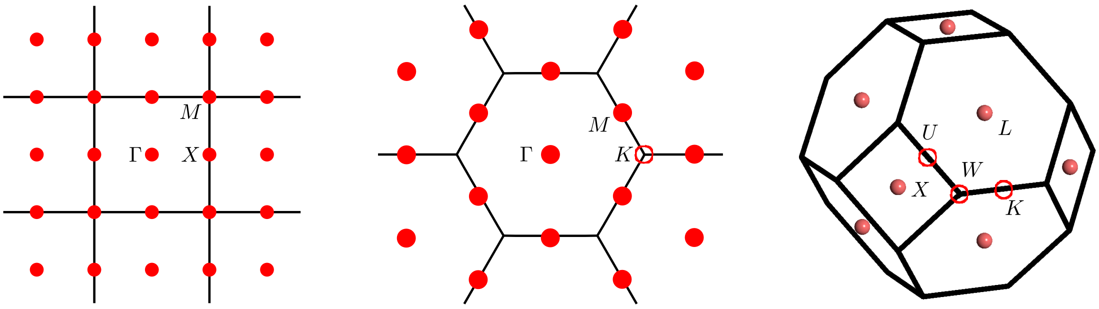

Figure 2.2.: *The two standard 2D lattices and the fcc–lattice with the respective sub–lattices of points at which the eigenvalue function is necessarily stationary due to time-reversal. The first Brillouin zones are outlined in black, Please note that the K-point in the triangular case (center) and the points K, W and U of the fcc–lattice (right) are not covered. In the illustration on the right hand side, the Γ point and all sub–lattice sites on the far side of the BZ were omitted to avoid confusion.*

coincide $\lambda(\mathbf{k}) = \lambda(-\mathbf{k})$. This means that the eigenvalue (not frequency) band cut along any straight line through Γ is an even function with respect to Γ and smooth due to holomorphy. Thus, solitary eigenvalue bands are stationary (minimum, maximum or saddle point) at Γ. Details about the periodicity were not required for this result, so it holds for any lattice.

In addition to the Γ–point, the eigenvalue bands are subject to inversion symmetry with respect to any point $\mathbf{k}_0$ that is equivalent to its image under inversion, i. e. , for which one reciprocal lattice vector $\mathbf{G}$ exists such that

$$\mathbf{k}_0 = \mathbf{G} - \mathbf{k}_0. \tag{2.16}$$

We find:

$$\begin{aligned}
\lambda(\mathbf{k}_0 + \delta\mathbf{k}) &= \lambda(-\mathbf{k}_0 - \delta\mathbf{k}) \quad \text{(inversion)} && (2.17)\\
&= \lambda(\mathbf{k}_0 - \mathbf{G} - \delta\mathbf{k}) \quad \text{(Eqn. (2.16))} && (2.18)\\
&= \lambda(\mathbf{k}_0 - \delta\mathbf{k}) \quad \text{(k–space periodicity)}. && (2.19)
\end{aligned}$$

Such points form a sub–lattice in reciprocal space which is defined by

$$\mathbf{k}_0 = \frac{1}{2}\mathbf{G}. \tag{2.20}$$

As in the case of Γ (see above), the gradient of the eigenvalue band vanishes and the band is stationary if solitary. Fig. 2.2 shows the sub–lattices for the two 2D cases with non–trivial lattice symmetry and for the 3D face–centered cubic (fcc) lattice. The only 2D high–symmetry point that is not contained in the sub–lattices is the K–point of the triangular lattice. In the fcc–case, only Γ, L and the X–point are covered.

Point group	Behavior of eigenvalue function	Example
$\mathcal{C}_s$	one directional derivative vanishes	many high symmetry lines
$\mathcal{C}_2$	stationary point	
$\mathcal{C}_3$	extremum	
$\mathcal{C}_4$	extremum	
$\mathcal{C}_6$	extremum	
$\mathcal{C}_{2v}$	stationary point	$X(\square)$, $M(\triangle)$
$\mathcal{C}_{3v}$	extremum	$K(\triangle)$
$\mathcal{C}_{4v}$	extremum	$\Gamma(\square)$, $M(\square)$
$\mathcal{C}_{6v}$	extremum	$\Gamma(\triangle)$

Table 2.1.: *Relation between the symmetry of some* **k***–points and the properties of an eigenvalue function that is solitary at these point for crystals with 2D lattice. In the cases denoted "extremum", a local extremum is guaranteed only if the second derivative of the function does not vanish. Otherwise, "higher order saddle–like points" are possible. The examples in the rightmost column refer to crystals that support the full point group of the lattice, i. e. where the basis does not break any symmetries. I denote the high–symmetry points by their canonical letters (cf. Ref. [20]) and the respective lattice (□ for square, △ for triangular lattice) in parentheses.*

So far, we have only shown the existence of stationary points rather than of band extrema. Furthermore, in many practically relevant triangular systems, the K–point partly defines the band gap although it is not in said sub–lattice. Both can be understood if the crystal is invariant under some point symmetries. Consider a point $\mathbf{k}_0$ which is not necessarily in the sub–lattice and let $\mathcal{P}^{(\mathbf{k}_0)}$[8] be its little cogroup in reciprocal space. In two dimensions, we find:

- In 2D, the two–fold rotation $\hat{C}_2$ is identical to inversion $\hat{I}$ and the results obtained above apply directly: The band is stationary at this point.
- A mirror plane imposes that the directional **k**–space derivative perpendicular to the mirror plane vanishes.
- A three–fold $\hat{C}_3$ or four–fold $\hat{C}_4$ rotation requires the first derivative to have three–fold or four–fold symmetry, respectively. Three or four directional derivatives can be identical only if the gradient vanishes. Furthermore, symmetry requires the second derivative to have three–fold (four–fold) symmetry, too. This excludes the possibility of a saddle point and predicts a local extremum unless the second derivative vanishes, too. However, there is no systematic reason for this to

---

[8] $\mathcal{P}^{(\mathbf{k}_0)}$ is the maximal subgroup of $\mathcal{P}^{(\Gamma)}$ whose elements map $\mathbf{k}_0$ on itself modulo a reciprocal lattice vector. $\mathcal{P}^{(\Gamma)}$ consists of all point operations in the photonic crystal's space group [20] plus **k**–space inversion due to time reversal symmetry.

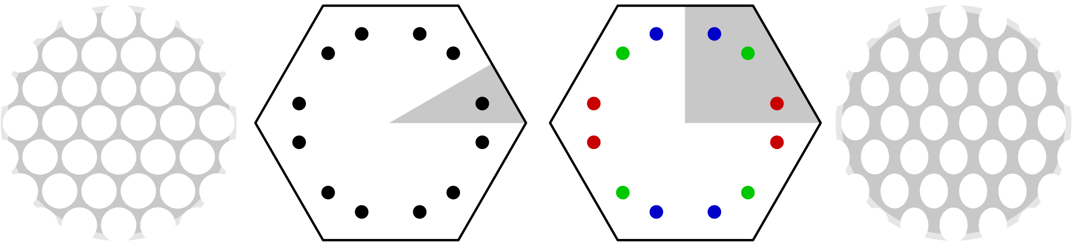

Figure 2.3.: *Sketch to illustrate the decomposition of a star in the first Brillouin zone into several "sub–stars" if the symmetry of a fully symmetric triangular system (left panel) is reduced to* $\mathcal{C}_{2v}$ *(right panel). In the fully symmetric case (left panel), all points of a star (black dots) are mapped onto each other by elements of the full point symmetry group. If symmetries are broken, these points form several stars with respect to the reduced point group. In the right panel, each color indicates one "sub–star". The irreducible wedges of the Brillouin zone are highlighted in gray.*

happen; therefore, such a feature in a band structure would be very sensitive to perturbations. Hence we might assume that indeed an extremum occurs.

From this, we can derive stationarity properties of points with respect to their **k**–space symmetry (cf. Tab. 2.1)

2.2.2. Consequences of Symmetry Breaking

As we have seen, symmetry can restrict time–reversal–related stationary points to extrema or can impose further ones. A very interesting problem is how the band structure and in particular the band–gap–defining band extrema change if a highly symmetric structure is modified such that symmetries are broken.

If the perturbation is small, there are three main consequences:

- The star of a **k**–point is divided into "sub–stars". Each element of a sub–star is mapped onto every other element by an operation of the reduced symmetry group but not on any element of one of the other sub–stars. This is illustrated in Fig. 2.3. As all points in a star are connected by symmetry operations, they are degenerate. Consequently, **k**–points that belong to different "sub–stars" are degenerate in the symmetric structure; this degeneracy is lifted in the perturbed system (cf. Fig. 2.4).

- If an extremum is required by symmetry in the original crystal but only a stationary point is guaranteed in the perturbed one, the extremum is distorted (e. g. the Γ–point in Fig. 2.5) and may become a saddle point if the perturbation is sufficiently strong. In the latter case, new extrema may be expected in the vicinity.

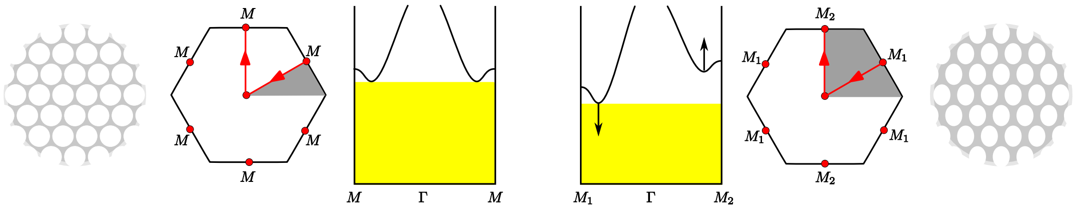

Figure 2.4.: *Sketch of the splitting of degeneracies on* $\mathbf{k}$*–space lines that are no longer connected by a symmetry operation due to a perturbation of the PhC as illustrated in Fig. 2.3. Here, I sketched the band structure along a path (red line with arrows) consisting of two lines of* $\overline{\Gamma M}$*–type in a triangular lattice with fully symmetric basis (plots on the left side). As every possible choice for the irreducible Brillouin zone (iBZ, for example the gray triangle) contains one such line and all those lines are equivalent, plotting a band structure only on the outline of the iBZ shows all extrema on this line. If the full symmetry is broken by the basis, a set of equivalent points can be divided into distinct classes (here:* M_1 *and* M_2*) and their degeneracy is lifted. In this case, care must be taken that members of all such classes are covered by the band structure path.*

.

- If a stationary point (or even an extremum) is required by symmetry in the original crystal but not in the perturbed one, the stationary point typically will move away from the high symmetry point in addition to being distorted (e. g. the K–point in Fig. 2.5).

To illustrate this, let us consider two examples: Macroporous silicon of cylindrical holes and as a—not only slightly—perturbed system the same lattice with rectangular holes. The corresponding band structures are depicted in Fig. 2.6.

In the band structure plot for the symmetric system, the band–free spectral range is defined by the first band at the K–point and by the second band at the M–point. Both points are expected to correspond to stationary points unless a degeneracy occurs. This can be verified by inspecting the band structure plot. At the high symmetry points, the band structure path abruptly changes its direction. Hence, the plot depicts the directional derivatives in two directions at these points. The first band has a saddle point at M and a local maximum at K. The second band has a local minimum at M, is degenerate at K and has another local minimum between Γ and K, because the symmetry of the $\overline{\Gamma K}$–line guarantees that the perpendicular directional derivative vanishes. If no further minima occur inside the irreducible part of the Brillouin zone (and they do not), the yellow–shaded area indeed is the band gap.

In the corresponding band structure for the system with reduced symmetry, there are now two classes of M–like points (M_1 and M_2) and two classes of K–like points (K_1 and K_2). Both M–like points are stationary because of time–reversal and the degeneracy

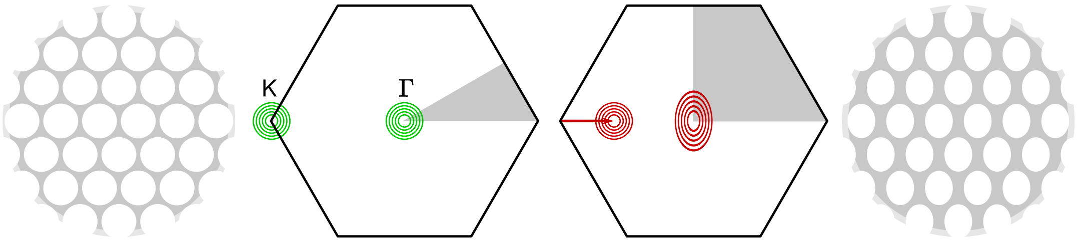

Figure 2.5.: *Sketch of how extrema can be modified depending on their position if the symmetry of a fully symmetric triangular system (left panel) is reduced to* $\mathcal{C}_{2v}$ *(right panel). The green and red lines represent isofrequency lines for a fully symmetric crystal and one with reduced symmetry, respectively. If the full symmetry requires an extremum at a certain point and the reduced symmetry guarantees the existence of a stationary point (e. g. the* Γ*–point in this sketch), the extremum remains at its position but may be deformed. If the perturbation is sufficiently strong, the former extremum may even change sign or become a saddle point (not shown). If the full symmetry requires an extremum at a high symmetry point but the reduced symmetry does not (e. g. the* K*–point in this sketch), the extremum furthermore tends to leave its spot in addition to a possible deformation (not shown).*

between them is lifted. However, the K–like points are no longer sufficiently symmetric to fix a stationary point. Instead, the extrema located at K in the fully symmetric case may shift anywhere. Furthermore, the K–like points remain degenerate because they are connected by reciprocal lattice translations.

In this plot, the band–free spectral range is bounded by the first band between M_1 and K_1 (or between M_1 and K_2) and by the second band between K_1 and Γ. The first band has saddle points at the M–like and local maxima somewhere near the K_1– and K_2–points. The second band has local minima at both M–like points where the frequency at M_1 is slightly higher than at M_2. Furthermore, it has a local minimum between Γ and K_1 because the symmetry of the $\overline{\Gamma K_1}$–line still guarantees that the perpendicular directional derivative vanishes. From this we may conclude that the upper edge of the band gap has been correctly identified in this plot whereas the maximum that defines the lower band gap edge moved away from K_2 to some general point inside the first Brillouin zone. This plot overestimates the size of the band gap.

2.2.3. Tight–Binding Model

In Sec. 2.2.1, we have seen that solitary bands are stationary at certain classes of high symmetry points. This defines a minimum number of stationary points for every isolated band. In the case of a band group, the situation is already a bit more complicated because degeneracies can occur at such high symmetry points and reduce the number

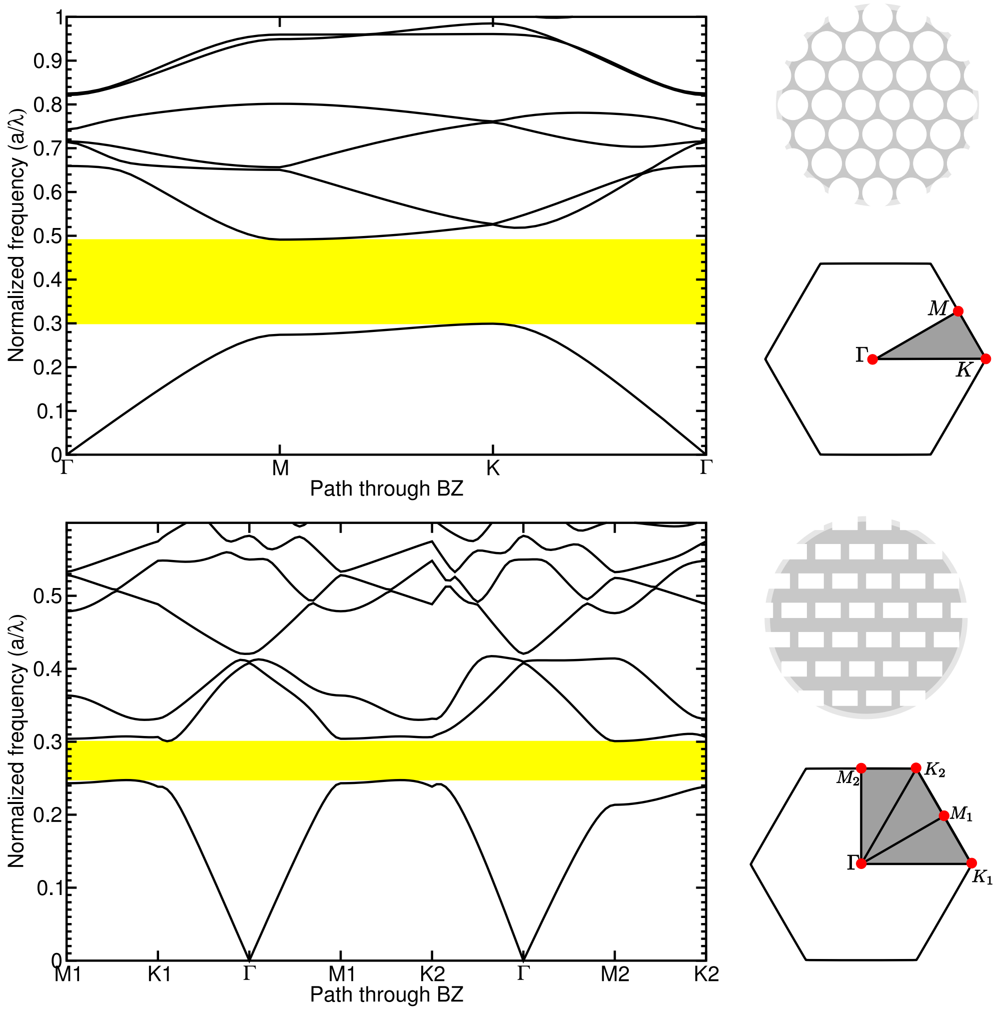

Figure 2.6.: *Band structure plots on the high-symmetry paths of the holohedry which are inequivalent due to symmetry breaking. Top: Triangular lattice of circular (radius* $0.45a$*) air pores which has full* $\mathcal{C}_{6v}$ *symmetry. Here, all solitary bands flatten at both the* M*– and the* K*–point; both points are stationary unless degenerate. Below: Triangular lattice of rectangular (dimensions* $0.5a \times 0.8a$*) air holes which only has* $\mathcal{C}_{2v}$ *symmetry. Here, the* M*–like points remain stationary unless degenerate whereas the* K*–like points loose this property. This is illustrated by the fact that the bands exhibit kinks if the band structure path changes its direction at these points. Both band structure calculations were performed with MPB at* 32×32 *discretization points per unit cell.*

of required extrema. On the other hand, bands in a band group often feature extrema on a line between a high symmetry point and Γ.

In this subsection, I try to estimate an upper bound for the number of stationary points. To this end, I approximate the PhC by a tight–binding model and derive some basic results for the simplest possible case: Pure next–neighbor coupling. This approximation is motivated by the fact that (at least in 2D) Wannier functions can be chosen to be tightly localized.

The wave equation has the form:

$$\hat{L}\mathbf{H}_n(\mathbf{k}) = \lambda_n(\mathbf{k})\mathbf{H}_n(\mathbf{k}), \tag{2.21}$$

where $\hat{L}$ is a linear operator and λ_n are its eigenvalues. First we formulate this in a Wannier basis (cf. Chap. 3), at this point still including the full long–range interaction. The coupling between the Wannier functions n and n' at the lattice sites $\mathbf{R}$ and $\mathbf{R}'$ is defined by the entries of the hopping matrix

$$\Lambda_{nn',\mathbf{R}-\mathbf{R}'} = \int \mathrm{d}r \Big(\mathbf{W}_{n\mathbf{R}}(\mathbf{r})\Big)^{*} \hat{L}\ \mathbf{W}_{n'\mathbf{R}'}(\mathbf{r}). \tag{2.22}$$

It is block–diagonal with respect to n and n' where each block corresponds to an isolated group of photonic bands. These blocks can be handled individually, hence we may assume that Λ only contains one such block.

The wave equation in Wannier formulation reads:

$$\det\left[\lambda - \sum_{\mathbf{R}} \exp\big(\mathrm{i}\mathbf{k}\mathbf{R}\big)\,\Lambda_{\mathbf{R}}\right] = 0, \tag{2.23}$$

where the summation runs over all real–space lattice translations. Thus, it is convenient to express $\mathbf{R}$ and $\mathbf{k}$ in the respective lattice coordinates

$$\begin{aligned} \mathbf{R} &= R_1\mathbf{a}_1 + R_2\mathbf{a}_2, \quad R_{1/2} \in \mathbb{Z}, & (2.24)\\ \mathbf{k} &= k_1\mathbf{b}_1 + k_2\mathbf{b}_2, \quad k_{1/2} \in (-1/2, 1/2], & (2.25) \end{aligned}$$

where $\mathbf{a}_{1/2}$ are the real–space and $\mathbf{b}_{1/2}$ the reciprocal lattice primitives.

We can thus expand the vectorial quantities:

$$\begin{aligned} 0 &= \det\left[\lambda - \sum_{R_1,R_2} \exp\big(2\mathrm{i}\pi k_1 R_1 + 2\mathrm{i}\pi k_2 R_2\big)\,\Lambda_{R_1R_2}\right] & (2.26)\\ &= \det\Bigg[\lambda - \underbrace{\sum_{R_1,R_2} z_1^{R_1} z_2^{R_2} \Lambda_{R_1R_2}}_{\mathcal{L}}\Bigg], & (2.27) \end{aligned}$$

where we used the substitution

$$z_{1,2} = \exp\Big(2\mathrm{i}\pi k_{1,2}\Big). \tag{2.28}$$

Each matrix element of $\mathcal{L}$ (being a Laurent series with respect to z_1 and z_2) is holomorphic on a ring in the complex z_1–plane and another ring in the z_2–plane. From the Hermiticity of our wave operator together with the time reversal symmetry follows

$$\Lambda_{nn',R_1R_2} = \Lambda_{nn',(-R_1)(-R_2)} \in \mathbb{R}. \tag{2.29}$$

If we assume z_2 to be fixed for a moment, we find:

$$\begin{aligned} & \sum_{R_1,R_2} z_1^{R_1} z_2^{R_2} \Lambda_{nn',R_1R_2} && (2.30)\\ = & \sum_{R_1,R_2} z_1^{-R_1} z_2^{-R_2} \Lambda_{nn',(-R_1)(-R_2)} && (2.31)\\ = & z_2^{-2R_2} \sum_{R_1,R_2} z_1^{-R_1} z_2^{R_2} \Lambda_{nn',R_1R_2}. && (2.32) \end{aligned}$$

The factor $z_2^{-2R_2}$ only introduces a phase. Thus, if a matrix element of $\mathcal{L}$ is holomorphic at z_1, it clearly is so at z_1^{-1} as well and consequently it is holomorphic on the unit circle of z_1 as long as the series converges at all. The same applies for z_2 if z_1 is kept fixed. This means that the coefficients of the characteristic polynomial are holomorphic functions of $\mathbf{k}$ throughout the whole Brillouin zone. Consequently, the eigenvalue bands are Riemann surfaces and thus locally holomorphic with respect to $\mathbf{k}$.

Stationary points in the eigenvalue bands are the points where both components of the $\mathbf{k}$–space gradient $\nabla_{\mathbf{k}}\lambda(\mathbf{k})$ vanish. Using the implicit function theorem, this is equivalent to finding the roots[9] of:

$$\begin{aligned} & \frac{\partial_{k_1} \det[\lambda - \mathcal{L}(k_1,k_2)]}{\partial_\lambda \det[\lambda - \mathcal{L}(k_1,k_2)]} = \frac{\partial_{k_2} \det[\lambda - \mathcal{L}(k_1,k_2)]}{\partial_\lambda \det[\lambda - \mathcal{L}(k_1,k_2)]} = 0, && (2.33)\\ \text{or} \quad & \frac{\partial_{z_1} \det[\lambda - \mathcal{L}(z_1,z_2)]}{\partial_\lambda \det[\lambda - \mathcal{L}(z_1,z_2)]} = \frac{\partial_{z_2} \det[\lambda - \mathcal{L}(z_1,z_2)]}{\partial_\lambda \det[\lambda - \mathcal{L}(z_1,z_2)]} = 0. && (2.34) \end{aligned}$$

Here, λ is the eigenvalue at the stationary point and we furthermore require $|z_{1,2}| = 1$. If the denominator vanishes, both numerators must do so as well. Otherwise, the band would have infinite group velocity, which would be unphysical.

Four–Neighbor Coupling

Now, we restrict ourselves to the tight–binding model in a square lattice, which means that all hopping matrices vanish except for $\Lambda_{10} = \Lambda_{-10}$, $\Lambda_{01} = \Lambda_{0-1}$ and Λ_{00}. The eigenvalue equation is:

$$\det[\lambda - \Lambda_{00} - (z_1 + z_1^{-1})\Lambda_{10} - (z_2 + z_2^{-1})\Lambda_{01}] = 0. \tag{2.35}$$

---

[9] The pair (z_1, z_2) that solves the equations for fixed λ

The derivatives in Eqn. (2.34) can be evaluated using Jacobi's formula

$$\partial_x \det[\mathcal{M}] = \mathrm{Tr}\left[\mathcal{M}^{\#}\partial_x\mathcal{M}\right], \tag{2.36}$$

where $\mathcal{M}^{\#}$ is the adjugate of the x–dependent matrix $\mathcal{M}(x)$. We find a condition that is equivalent to Eqn. (2.34):

$$\begin{aligned}
\partial_{z_1}\det[\lambda-\mathcal{L}] &= \mathrm{Tr}\left[(\lambda-\mathcal{L})^{\#}\cdot(1-z_1^{-2})\Lambda_{10}\right] \\
&= (1-z_1^{-2})\mathrm{Tr}\left[(\lambda-\mathcal{L})^{\#}\cdot\Lambda_{10}\right]=0, \qquad (2.37)\\
\partial_{z_2}\det[\lambda-\mathcal{L}] &= \mathrm{Tr}\left[(\lambda-\mathcal{L})^{\#}\cdot(1-z_2^{-2})\Lambda_{01}\right] \\
&= (1-z_2^{-2})\mathrm{Tr}\left[(\lambda-\mathcal{L})^{\#}\cdot\Lambda_{01}\right]=0, \qquad (2.38)\\
\partial_{\lambda}\det[\lambda-\mathcal{L}] &= \mathrm{Tr}\left[(\lambda-\mathcal{L})^{\#}\right]\neq 0. \qquad (2.39)
\end{aligned}$$

Let us first discuss Eqn. (2.37). It obviously has the trivial zeros $z_1=\pm1$, which create ridges or valleys along lines of high symmetry. For a band group, however, additional ridges can occur if the trace becomes zero. The same holds for Eqn. (2.38).

If Λ_{10} is invertible, we find

$$\mathrm{Tr}\left[(\lambda-\mathcal{L})^{\#}\cdot\Lambda_{10}\right] = \mathrm{Tr}\left(\left[\Lambda_{10}^{-1}(\lambda-\mathcal{L})\right]^{\#}\right), \tag{2.40}$$

$$= \mathrm{Tr}\left(\left[\Lambda_{10}^{-1}(\lambda-\Lambda_{00}+(z_1+z_1^{-1})\Lambda_{10}+(z_2+z_2^{-1})\Lambda_{01})\right]^{\#}\right). \tag{2.41}$$

Each of the adjugate's elements is the determinant of a $(N-1)\times(N-1)$ matrix, where N is the number of rows of $\mathcal{L}$. Consequently, Eqn. (2.41) can be recast to a polynomial in z_1 and z_2 of degree $2N-2$, whose zeros come in complex conjugate pairs, because of $z_{1/2}^{-1}=z_{1/2}^{*}$. The finite degree of the polynomial imposes an upper bound of $2N$ for the total number of ridges (including the trivial zeros) for each band in an entangled tight–binding group. From that we may deduce that for a single tight–binding band, only the trivial zeros exist[10].

Extrema occur at the intersections of two ridges or of two valleys. The trivial zeros $z_1=\pm1$ and $z_2=\pm1$ ($k_{1/2}=n/2,\ n\in\mathbb{Z}$) define the points which we already found to be necessarily stationary if solitary (cf. Sec. 2.2.1). A nontrivial zero in one directional derivative and a trivial one in the other derivative creates a stationary point on a high symmetry line between Γ and one of the necessarily stationary points (e. g. X and M in 2D and X or L in the fcc–lattice). Finally, if a nontrivial zero occurs in both directional derivatives, this corresponds to a stationary point on some other high symmetry line or at a general point of the first Brillouin zone. This means that even in the very simple case of multiple–band tight–binding coupling additional extrema can appear essentially anywhere in the Brillouin zone.

---

[10] For a single band, the trace simply becomes a non–zero constant.

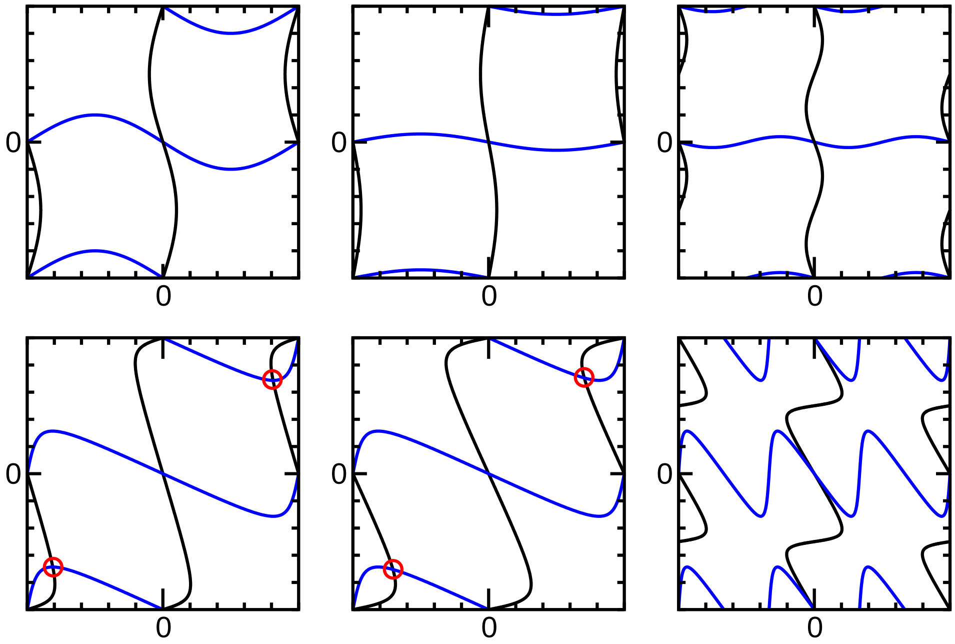

Figure 2.7.: *Qualitative sketch of the points of vanishing directional derivatives in the presence of a single long–range interaction term. Each panel depicts a reciprocal unit cell with k_1 and k_2 varying along the horizontal and the vertical axis, respectively. Black lines indicate wave vectors that fulfill Eqn. (2.44) (first component of gradient vanishes) and blue lines those that obey Eqn. (2.45). The left column illustrates the case $|\Lambda_{10}| \neq |\Lambda_{01}|$ with a weak (top row) and a strong (bottom row) perturbation term Λ_{11}. A distorted triangular lattice of circles would be an example for this situation. The center column shows similar plots for $\Lambda_{10} = \Lambda_{01}$. This corresponds to a triangular lattice of ellipses that are oriented along a direction of high symmetry. Finally, the rightmost column depicts the situation for a further distance coupling Λ_{32} again with $\Lambda_{10} \neq \Lambda_{01}$. Stationary points are highlighted by red circles.*

Long–Range Interaction

Finally, coupling between lattice sites that are neither of the four formal next neighbors is the second possible source of additional extrema. A crystal with the full symmetry of the triangular lattice is an example for this; each lattice site couples equally strong to its six neighbors. Concurrently, the occurrence of stationary points at the K–points is imposed by symmetry (cf. Sec. 2.2.1).

The question arises how great a long–range coupling term must be to create additional extrema. To simplify matters, I would like restrict myself to the single band case. Consequently, the coupling matrices Λ_{nm} become scalars. Let us consider a single additional coupling term between lattice sites that are separated by some lattice vector $\mathbf{R} = n\mathbf{a}_1 + m\mathbf{a}_2$. The resulting scalar discretized system operator reads:

$$\begin{aligned} \mathcal{L} &= \Lambda_{00} + (z_1 + z_1^{-1})\Lambda_{10} + (z_2 + z_2^{-1})\Lambda_{01} + (z_1^n z_2^m + z_1^{-n} z_2^{-m})\Lambda_{nm} \qquad (2.42)\\ &= \Lambda_{00} + 2\Lambda_{10}\cos(2\pi k_1) + 2\Lambda_{01}\cos(2\pi k_2) + 2\Lambda_{nm}\cos(2\pi n k_1 + 2\pi m k_2). \qquad (2.43) \end{aligned}$$

Its derivatives with respect to k_1 and k_2 are:

$$\begin{aligned} \partial_{k_1}[\lambda - \mathcal{L}] &= 4\pi\Lambda_{10}\sin(2\pi k_1) + 4\pi n\Lambda_{nm}\sin(2\pi n k_1 + 2\pi m k_2) = 0, \qquad (2.44)\\ \partial_{k_2}[\lambda - \mathcal{L}] &= 4\pi\Lambda_{01}\sin(2\pi k_2) + 4\pi m\Lambda_{nm}\sin(2\pi n k_1 + 2\pi m k_2) = 0. \qquad (2.45) \end{aligned}$$

First, we note that the trivial zeros $2k_{1/2} \in \mathbb{Z}$ are still present. Furthermore, this set of equations can feature additional zeros depending on the coupling constants. A closed form for these solutions in terms of inverse trigonometric functions is easily found but does not provide too much insight. Instead, we can sketch the **k**–space lines on which Eqn. (2.44) and Eqn. (2.45) are fulfilled for some typical cases. Examples are shown in Fig. 2.7. The figure shows that non–trivial intersections of the lines of vanishing directional derivatives occur only if the perturbation $|\Lambda_{nm}|$ exceeds a certain threshold. Furthermore, it suggests that a necessary condition for the appearance of non–trivial zeros is that the slope $(\mathrm{d}k_2)/(\mathrm{d}k_1)$ of the blue curve is somewhere greater than the black one's. This is tantamount to the condition

$$\left(\frac{\mathrm{d}k_1}{\mathrm{d}k_2}\right)^{(\mathrm{I})} \cdot \left(\frac{\mathrm{d}k_2}{\mathrm{d}k_1}\right)^{(\mathrm{II})} \geq 1, \qquad (2.46)$$

where the superscripts (I) and (II) indicate that within the parenthesis k_1 and k_2 fulfill Eqn. (2.44) or Eqn. (2.45), respectively. The implicit function theorem states that

$$\begin{aligned} \left(\frac{\mathrm{d}k_1}{\mathrm{d}k_2}\right)^{(\mathrm{I})} &= \frac{\partial_{k_2}\partial_{k_1}[\lambda - \mathcal{L}]}{\partial_{k_1}^2[\lambda - \mathcal{L}]} \qquad (2.47)\\ &= \frac{m}{n}\left[1 - \Lambda_{10}\Big(\Lambda_{10} + \frac{n^2\Lambda_{nm}\cos(2\pi n k_1 + 2\pi m k_2)}{\cos(2\pi k_1)}\Big)^{-1}\right], \qquad (2.48)\\ \left(\frac{\mathrm{d}k_2}{\mathrm{d}k_1}\right)^{(\mathrm{II})} &= \frac{\partial_{k_1}\partial_{k_2}[\lambda - \mathcal{L}]}{\partial_{k_2}^2[\lambda - \mathcal{L}]} \qquad (2.49)\\ &= \frac{n}{m}\left[1 - \Lambda_{01}\Big(\Lambda_{01} + \frac{m^2\Lambda_{nm}\cos(2\pi n k_1 + 2\pi m k_2)}{\cos(2\pi k_1)}\Big)^{-1}\right]. \qquad (2.50) \end{aligned}$$

Obviously, the prefactors m/n and n/m cancel if these expressions are inserted in Eqn. (2.46). The main effect of n and m is that they scale the long–range coupling term Λ_{nm}. Fig. 2.7 shows that the points of greatest slope are at the zeros of the sine functions in Eqn. (2.44) and Eqn. (2.45). From this follows that the critical threshold for a long–range coupling term to cause additional extrema decays slowly; roughly according to $\Lambda_{nm}^{(\text{crit.})} \sim \min\{n^{-2}, m^{-2}\}$.

We conclude that the band structure is to a certain extent robust with respect to the occurrence of additional stationary points due to long–range interaction. However, we can see in Fig. 2.7 that they in principle may appear anywhere in the first Brillouin zone unless the hopping coefficients fulfill fixed relations, e. g. due to symmetry.

If the crystal can be regarded as a composition of coupled resonators, one may assume that the actual long–range coupling decays exponentially. As a result, only very few terms such as Λ_{11} in a triangular lattice can overcome the threshold for creating stationary points, the K in systems with three–fold symmetry being one notable example.

2.2.4. Rules of Thumb

Here, I would like to condense the results obtained above to a few rules of thumb for deducing the band gap of a new PhC from a band structure plot. They have proven to be useful in my calculations but are most likely incomplete and may occasionally fail.

First, the symmetry of the structure has to be determined. If the structure does not feature all symmetries of the lattice, there might exist a closely related system with full symmetry[11] which has already been well studied. This should already give a good first impression of the band structure of the distorted system.

One should bear in mind:

1. The target band structure should be plotted at least on the outline of the irreducible Brillouin zone of the *reduced* symmetry. If the related fully symmetric system features extrema on high symmetry lines, all such lines inside the irreducible wedge of the reduced symmetry should be included (as in Fig. 2.6). One should try to relocate all extrema of the related symmetric band structure if they are close to the band gap.

2. In square or rectangular lattices, all high symmetry points are stationary due to time reversal. Extrema at such points do not leave their spots if symmetries are broken. Extrema on lines connecting high symmetry points do, though.

3. In the triangular lattice, all high symmetry points are stationary as long as the basis supports a three–fold axis. If this symmetry is broken, extrema at Γ and M do not move whereas those at K do. In general, the latter extrema will not stay on a connecting line.

4. If the band structure plotted on a path along lines of high symmetry exhibits kinks where the path abruptly changes direction (typically at high symmetry points), this is a clear sign of a stationary point next to the path. The sharpness of the

---

[11] For instance, replace the basis by a set of spheres.

kink roughly indicates how far away the stationary point is and how severely the band gap size may be overestimated by the present band structure.

5. An isolated tight–binding band has stationary points only where required by time reversal. Long range interaction must exceed a certain threshold to create additional extrema, which first appear close to points that are stationary due to time reversal. The situation is similar for interband coupling, which in highly symmetric systems seems to be the main reason for extrema at positions with low symmetry[12]. The number of stationary points in an entangled group of N tight–binding bands is limited to $2N$.

There are important exceptions from the fourth rule: In the fcc–lattice (even for a fully symmetric basis), the K–point and the U–point[13] are not sufficiently symmetric to fix stationary points. Thus, a kink at these points is not uncommon. Usually, there is a nearby extremum on the line $\overline{\Gamma K}$; this point is stationary in sufficiently symmetric systems.

---

[12] Examples are the second band on the $\overline{\Gamma M}$–line in square arrangements of posts or holes and the second band on the $\overline{\Gamma K}$–line in similar triangular systems.

[13] In fact, those points are equivalent

2.3. On–Shell Band Structure — Task and Solution Properties

In Sec. 2.1, we already discussed the conventional band structure approach where the wave vector is given as a parameter and the corresponding allowed frequencies are calculated. However, it is equally possible to follow the inverse approach where the frequency is a fixed parameter and possible wave vectors are determined. This is the so–called on–shell band structure problem.

It is a valuable tool for a variety of problems such as density of states calculations (cf. Sec. 2.3.1 and Chap. 5) and the refraction problem at a photonic crystal surface (cf. Chap. 4). Another possible application for this technique is the investigation of thermal emission from a photonic crystal, which was done by Christian Schuler as the topic of his diploma thesis in our group [6, 25], which was partly supervised by me.

2.3.1. Density of States

The total density of states (DOS) of a crystal is the "number" of wave vectors that correspond to eigenmodes with a given frequency ω. As $\mathbf{k}$ is a continuous variable, this translates to an integral over the first BZ:

$$N(\omega) = \sum_n \int_{\mathrm{BZ}} \mathrm{d}k \; \delta(\omega_{n\mathbf{k}} - \omega). \tag{2.51}$$

Here, $\omega_{n\mathbf{k}}$ is the eigenfrequency of the n–th band at $\mathbf{k}$ and is regarded as a function of $\mathbf{k}$. The density of states simply is the total area of the isofrequency surfaces (IFS), i. e. in each band the set of $\mathbf{k}$–vectors that corresponds to Bloch modes with said frequency:

$$N(\omega) = \sum_n \int_{\mathrm{IFS}_n(\omega)} \mathrm{d}k, \quad \mathrm{IFS}_n(\omega) = \{\mathbf{k} : \omega_{n\mathbf{k}} = \omega\}. \tag{2.52}$$

However, in many cases such as calculating the electromagnetic loss channels of a localized object, this quantity is not very meaningful. Instead, we need the "number" of $\mathbf{k}$–vectors that can be coupled to. This is achieved by weighting each state with the probability of the object to interact with this particular mode.

If the coupling is mediated by a point–like electric dipole with dipole moment $\mathbf{d}$ situated at $\mathbf{r}$, the corresponding electrically projected local density of states (LDOS) reads:

$$\begin{aligned} N_{\mathrm{p,el}}(\omega, \mathbf{r}) &= \sum_n \int_{\mathrm{BZ}} \mathrm{d}k \delta(\omega_{n\mathbf{k}} - \omega) \left|\mathbf{d} \cdot \mathbf{E}_{n\mathbf{k}}(\mathbf{r})\right|^2, && (2.53) \\ &= \sum_n \int_{\mathrm{IFS}_n(\omega)} \mathrm{d}k \left|\mathbf{d} \cdot \mathbf{E}_{n\mathbf{k}}(\mathbf{r})\right|^2. && (2.54) \end{aligned}$$

Likewise, we can also define the magnetically projected local density of states for a magnetic dipole $\boldsymbol{\mu}$ as well:

$$\begin{aligned} N_{\mathrm{p,mag}}(\omega, \mathbf{r}) &= \sum_n \int_{\mathrm{BZ}} \mathrm{d}k \delta(\omega_{n\mathbf{k}} - \omega) \left| \boldsymbol{\mu} \cdot \mathbf{H}_{n\mathbf{k}}(\mathbf{r}) \right|^2, & (2.55) \\ &= \sum_n \int_{\mathrm{IFS}_n(\omega)} \mathrm{d}k \left| \boldsymbol{\mu} \cdot \mathbf{H}_{n\mathbf{k}}(\mathbf{r}) \right|^2. & (2.56) \end{aligned}$$

All three quantities can be calculated for a specific frequency ω by numerically integrating the respective integrands over the IFS. This naturally asks for a method to solve for the allowed wave vectors with the frequency being a parameter, hence, an on–shell band structure method.

2.3.2. Plane Wave Expansion

Analogously to the case of conventional band structure calculations, the wave equation can be discretized by expanding it in the basis of plane waves:

$$\sum_{\mathbf{G}'} \eta_{\mathbf{G}-\mathbf{G}'} (\mathbf{k} + \mathbf{G}) \times (\mathbf{k} + \mathbf{G}') \times \mathbf{h}_{n\mathbf{k}}^{(\mathbf{G}')} = \frac{\omega^2}{c^2} \mathbf{h}_{n\mathbf{k}}^{\mathbf{G}}. \tag{2.57}$$

This time the unknown for which we would like to solve the equation is the wave vector rather than the frequency. To this end, we decompose the wave vector into a fixed $\mathbf{k}_0$ and a variable contribution $\lambda \hat{e}_k$:

$$\mathbf{k} = \mathbf{k}_0 + \lambda \hat{e}_k. \tag{2.58}$$

$\hat{e}_k$ is the search direction in reciprocal space and λ is the eigenvalue of a (nonlinear) matrix eigenproblem:

$$\begin{aligned} \frac{\omega^2}{c^2} \mathbf{h}_{n\mathbf{k}}^{\mathbf{G}} - \sum_{\mathbf{G}'} \eta_{\mathbf{G}-\mathbf{G}'} \Bigg(& \lambda^2 \hat{e}_k \times \hat{e}_k \times \mathbf{h}_{n\mathbf{k}}^{(\mathbf{G}')} \\ & + \lambda \Big[(\mathbf{k}_0 + \mathbf{G}) \times \hat{e}_k \times \mathbf{h}_{n\mathbf{k}}^{(\mathbf{G}')} + \hat{e}_k \times (\mathbf{k}_0 + \mathbf{G}') \times \mathbf{h}_{n\mathbf{k}}^{(\mathbf{G}')} \Big] \\ & + (\mathbf{k}_0 + \mathbf{G}) \times (\mathbf{k}_0 + \mathbf{G}') \times \mathbf{h}_{n\mathbf{k}}^{(\mathbf{G}')} \Bigg) = 0. \end{aligned} \tag{2.59}$$

This equation is quadratic with respect to the eigenvalue λ. It can be transformed to a conventional (linear) eigenvalue problem with twice the number of unknowns.

For conventional band structures, we would now express the vectorial Fourier coefficients $\mathbf{h}_{n\mathbf{k}}^{(\mathbf{G})}$ in a $\mathbf{k}$–dependent coordinate system that consisted of a unit vector $\hat{e}_{||}$ parallel to $(\mathbf{k} + \mathbf{G})$ and two unit vectors $\hat{e}_{\perp,1/2}$ perpendicular to $(\mathbf{k} + \mathbf{G})$. In an on–shell band structure problem, the wave vector is unknown. Thus, this decomposition is not possible for three–dimensional systems and in this case the unphysical modes cannot be excluded a priori. As a consequence, the resulting matrix problem is larger by a factor

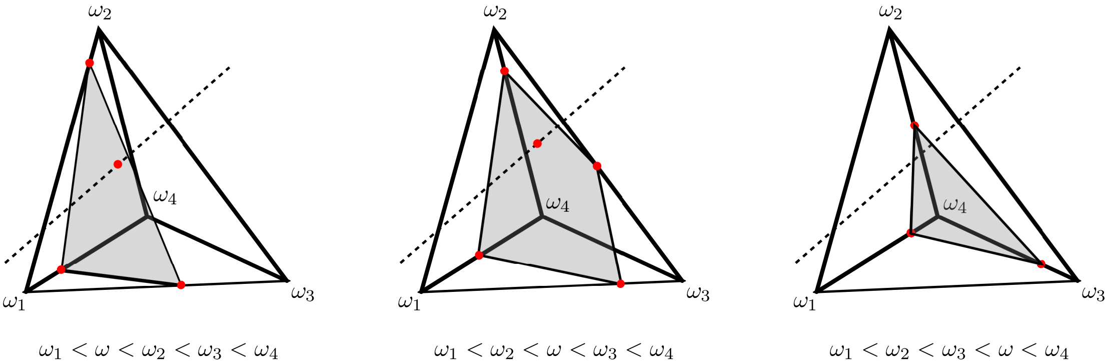

Figure 2.8.: *Sketch of the linearly approximated isofrequency contour (gray shaded polygons) in a three–dimensional simplex with four different frequencies associated with its corners. This sketch illustrates the three general, non–trivial possibilities for the operational frequency* ω *with respect to the corner frequencies. The vertices (red dots) of the isofrequency polygon divide the lines between* ω_i *and* $\omega_j > \omega > \omega_i$ *according to the ratio* $\frac{\omega_j-\omega}{\omega-\omega_i}$. *The black dotted lines indicate a line such as Eqn. (2.58), which should be intersected with the isofrequency contour*

of three (a factor of two due to the linearization and another factor of 1.5 due to the divergence condition) and accordingly solving it numerically is slower by a factor of 27 compared to a conventional plane wave calculation of comparable accuracy.

It should be noted that the unphysical solutions can be excluded for two–dimensional systems because here the polarizations decouple. However, the equations are messy [6] and the numerical effort still exceeds that for a conventional band structure calculation by a factor of eight. This demonstrates that on–shell solvers based on a plane wave expansion tend to be inefficient and other methods should be preferred.

2.3.3. Linear k–Space Interpolation

A relatively easy and fairly time–efficient way to calculate on–shell band structures is to perform conventional band structure calculations on a discretization (e.g. on an equidistant mesh) of the first Brillouin zone and to interpolate linearly between the resulting frequencies and Bloch modes. This approach is well–known in solid state theory [34, 35] and has been adapted to photonic crystals over a decade ago [32].

First, the Brillouin zone is tessellated with simplices (triangles in two and tetrahedra in three dimensions) whose vertices are the discretization points for the conventional band structure calculation. On each simplex, the frequency is linearly interpolated between the points of support. Consequently, isofrequency contours inside such a simplex form straight lines or flat polygons in two or three dimensions, respectively (cf. Fig. 2.8). This has the advantage that there are closed solutions for the area of the isofrequency contour and for the intersection between it and a line in reciprocal space such as Eqn. (2.58).

Thus, the density of states and the **k**–vectors associated with a frequency can be easily calculated by summing over the contributions of all simplices. It should be noted that the isofrequency tiles from all simplices indeed form a closed surface because each two neighboring simplices have the same linear interpolation on the common simplex face.

If the eigenmode for the determined wave vector **k** is required as well, the lattice–periodic parts of the Bloch modes at the points of support are linearly interpolated inside the simplex and afterward multiplied with the Bloch phase $\exp\left(\mathrm{i}\mathbf{kr}\right)$. Finally, the integrals Eqn. (2.53) and Eqn. (2.55) can also be analytically evaluated for Bloch functions whose periodic parts vary linearly with respect to **k** [32].

A practical drawback of this method is the first order convergence with respect to the **k**–space discretization due to the linear approximation. It could be improved by higher order interpolation polynomials. However, there would be no closed solutions (at least for sufficiently high polynomial orders) to the integration of an interpolated function over the isofrequency contour or the intersection with a **k**–space line. On the other hand, the linear interpolation provides reasonable accuracy of a few percent on a modern machine with enough memory. Finally, the most time–consuming part of this approach are the conventional band structure calculations on the points of support, which can be trivially parallelized on a computer cluster.

3

Chapter 3.

Wannier Functions

Since their introduction, Wannier functions (WF) have been of great value in the theoretical investigation (in particular of defect states) of electronic systems such as semiconductors. As we have seen in Sec. 2.2.3, they can be useful to some extent for understanding the properties of photonic crystals as well. Besides that, they can be employed as a basis for numeric simulations (cf. Sec. 3.5).

After stating the definition of Wannier functions in Sec. 3.1, I introduce the most popular approach to numerically create well–localized WFs in Sec. 3.2 and present the results of my efforts to improve their quality and speed up the generation in Sec. 3.3. Afterward, I discuss problems I encountered in the attempt to create favorable Wannier functions for three–dimensional photonic crystals in Sec. 3.4. This chapter is concluded with Sec. 3.5 about the application of photonic Wannier functions for the simulation of PhC–based photonic circuits. Here, the emphasis is on the problem of absorbing boundaries for PhC wave guides.

3.1. Definition

In this section, I provide the definition of Wannier functions and discuss some basic properties. Although they are not a feature of the band structure, WFs are closely related and have proven to be very useful for investigating it. Besides that, they also form an advantageous expansion basis for numerically investigating perturbed photonic crystals such as wave guiding structures (cf. Sec. 1.2.4) or disordered systems [15].

Wannier functions are the lattice–Fourier transform of the photonic crystal's Bloch functions:

$$\mathbf{W}_{n\mathbf{R}}(\mathbf{r}) = \frac{1}{V_{\mathrm{BZ}}} \int_{\mathrm{BZ}} \mathrm{d}k \, \exp\big(-\mathrm{i}\mathbf{k}\mathbf{R}\big)\, \mathbf{H}_{n\mathbf{k}}(\mathbf{r}). \tag{3.1}$$

In this formulation, Wannier functions are derived from the magnetic components of the Bloch modes. An analogous WF set in an electric field formulation can be constructed as well and it is the canonical choice to describe 2D photonic crystals when the light field can be completely described by a single electric field component [29]. Concerning the generation, there is no difference compared to the magnetic formulation, except for replacing the letter **H** with **E**.

For nearly all practically relevant photonic crystals in more than one dimension[1], the Bloch modes are not known analytically and have to be calculated numerically using band structure solvers (cf. Sec. 2.1). Consequently, the WFs have to be numerically constructed as well and the continuous integration in Eqn. (3.1) must be replaced with a sum over a discrete mesh $\mathcal{L}_{\mathbf{k}}$ of $\mathbf{k}$–points with N_k elements:

$$\mathbf{W}_{n\mathbf{R}}(\mathbf{r}) = \frac{1}{N_k} \sum_{\mathbf{k}\in\mathcal{L}_{\mathbf{k}}} \exp\left(-\mathrm{i}\mathbf{k}\mathbf{R}\right) \mathbf{H}_{n\mathbf{k}}(\mathbf{r}). \tag{3.2}$$

In our group, we always use a Monkhorst–Pack mesh [36].

As we have seen in Sec. 2.1, the frequency bands of a photonic crystal undergo several crossings and anticrossings throughout the first Brillouin zone. Each of these redistributes the modal patterns across the involved bands and Eqn. (3.2) averages over all such patterns. Consequently, the Bloch modes (as functions of $\mathbf{k}$ rather than of $\mathbf{r}$) are subject to abrupt changes, which leads to poor localization of the Fourier transform. Finally, exponential localization of Wannier functions that stem from a closed group of bands can only be obtained if the Bloch modes are hybridized in a specific fashion [22].

These three requirements can be met by replacing the Bloch modes in Eqn. (3.1) and Eqn. (3.2) with generalized Bloch modes that are obtained by applying a $\mathbf{k}$–dependent unitary mixing to the Bloch base:

$$\widetilde{\mathbf{H}}_{n\mathbf{k}}(\mathbf{r}) = \sum_m \mathcal{U}^{(\mathbf{k})}_{mn} \mathbf{H}_{m\mathbf{k}}(\mathbf{r}). \tag{3.3}$$

It is advisable to think of $\mathcal{U}^{(\mathbf{k})}$ as a matrix–valued function $\mathcal{U}(\mathbf{k})$.

Using this transformation, we arrive at an equation that defines a set of generalized Wannier functions

$$\mathbf{W}_{n\mathbf{R}}(\mathbf{r}) = \frac{1}{N_k} \sum_{\mathbf{k}\in\mathcal{L}_{\mathbf{k}}} \exp\left(-\mathrm{i}\mathbf{k}\mathbf{R}\right) \sum_m \widetilde{\mathbf{H}}_{m\mathbf{k}}(\mathbf{r}). \tag{3.4}$$

Just like the Wannier functions defined in Eqn. (3.2), this set of functions spans the same space as the Bloch modes from which it was created.

In numerical simulations (Sec. 3.5), the differences between proper WFs and generalized ones are of no practical relevance and we will not distinguish between them. Two of the most convenient properties of the Wannier basis are the translational symmetry and the orthogonality. The discrete translational invariance of the photonic crystal is inherited by the Wannier basis. The manifestation of this symmetry is that each Wannier function is associated with one lattice site $\mathbf{R}$ and the WF for the same band at a different site $\mathbf{R}'$ is identical except for a lattice translation:

$$\mathbf{W}_{n\mathbf{R}'}(\mathbf{r}) = \mathbf{W}_{n\mathbf{R}}(\mathbf{r} - \mathbf{R}' + \mathbf{R}). \tag{3.5}$$

---

[1] All relevant systems that I am aware of.

As a consequence, it is sufficient to construct the Wannier basis at the origin $\mathbf{R} = \mathbf{0}$.

The orthogonality of the Wannier basis, which can be specialized to an orthonormality

$$\langle\, \mathbf{W}_{n\mathbf{R}} \,|\, \mathbf{W}_{n'\mathbf{R}'} \,\rangle = \delta_{nn'}\delta_{\mathbf{R}\mathbf{R}'} \tag{3.6}$$

by properly normalizing the WFs, stems from the orthogonality of the Bloch basis. This, in turn, is due to the magnetic wave operator being Hermitian[2]. Here, I used a space–saving bra–ket notation:

$$\langle\, \mathbf{W}_{n\mathbf{R}} \,|\hat{A}|\, \mathbf{W}_{n'\mathbf{R}'} \,\rangle = \int \mathrm{d}r \mathbf{W}^{*}_{n\mathbf{R}}(\mathbf{r})\hat{A}(\mathbf{r})\mathbf{W}_{n'\mathbf{R}'}(\mathbf{r}). \tag{3.8}$$

3.2. Review of the Basic Generation

This section provides a brief introduction to the basic approach for creating Wannier functions that are well–suited for the numerical simulation techniques described in Sec. 3.5. The numerical construction of well–localized Wannier functions was a tough problem until 1997, when N. Marzari and coworkers published their seminal work [37] on efficiently optimizing the localization of WFs that are constructed from an isolated group of bands. It was extended to entangled band groups via an iterative subspace optimization developed by I. Souza and coworkers [38]. These advances in the realm of solid state theory laid the foundation for the development of the photonic Wannier functions method [39–41].

First, I briefly review the methods developed by N. Marzari and I. Souza in Sec. 3.2.1 and Sec. 3.2.2, respectively. Then, I comment on proper starting points for the minimization procedure. Finally, (Sec. 3.2.4), I describe the bottom–up approach for constructing photonic WFs, which was developed by my predecessor M. Schillinger [29].

3.2.1. Marzari–Vanderbilt Method for Isolated Band Groups

Wannier functions that are created from numerically obtained Bloch modes according to Eqn. (3.2) suffer from very poor localization. This renders them useless for any numerical treatment. As already mentioned in Sec. 3.1, acceptable localization can be obtained by switching to generalized Wannier functions

$$\mathbf{W}_{n\mathbf{R}}(\mathbf{r}) = \frac{1}{N_k} \sum_{\mathbf{k}\in\mathcal{L}_\mathbf{k}} \exp\left(-\mathrm{i}\mathbf{k}\mathbf{R}\right) \sum_{m} \mathcal{U}^{(\mathbf{k})}_{mn} \mathbf{H}_{m\mathbf{k}}(\mathbf{r}), \tag{3.9}$$

where the unitary mixing matrices $\mathcal{U}^{(\mathbf{k})}_{mn}$ have to be calculated numerically to fit the (numerically constructed) Bloch modes. To this end, N. Marzari proposed an iterative

---

[2] At this point, there actually is a slight difference between the electric and the magnetic field formulation: In the electric field formulation, the orthonormality relation contains the dielectric distribution as a kernel

$$\langle\, \mathbf{W}_{n\mathbf{R}} \,|\epsilon|\, \mathbf{W}_{n'\mathbf{R}'} \,\rangle = \delta_{nn'}. \tag{3.7}$$

approach which creates a sequence of mixings $\mathcal{U}^{(\mathbf{k}),(i)}$, where i enumerates the iteration steps. These mixings correspond to gradually better WFs.

To be more specific, the construction of Wannier functions is considered as a minimization problem where the objective function (i. e. the function that shall be minimized) is identified with the accumulated spread of the WFs

$$\Omega = \sum_n \left(\langle\, \mathbf{W}_{n\mathbf{0}}\, | \mathbf{r} |\, \mathbf{W}_{n\mathbf{0}}\, \rangle \right)^2 - \langle\, \mathbf{W}_{n\mathbf{0}}\, | r^2 |\, \mathbf{W}_{n\mathbf{0}}\, \rangle. \tag{3.10}$$

The degrees of freedom in this optimization problem are the unitary matrices. The naive notation in Eqn. (3.10) suggests to construct the Wannier basis and calculate their spreads in each iteration. This can be avoided by inserting Eqn. (3.9) in Eqn. (3.10) and exploiting the relation of the position operator to the $\mathbf{k}$–space gradient:

$$\hat{\mathbf{r}} \to \mathrm{i}\nabla_{\mathbf{k}}. \tag{3.11}$$

Consequently, the objective can be reformulated in terms of the Bloch modes (more specifically: of their lattice–periodic parts) and—with the help of $\mathbf{k}\cdot\mathbf{p}$ perturbation theory[3]—further approximated by first order finite differences:

$$\langle\, \mathbf{h}_{n\mathbf{k}}\, | \nabla_{\mathbf{k}} |\, \mathbf{h}_{m\mathbf{k}}\, \rangle \approx \sum_{\mathbf{b}} \mathbf{b} w_b \underbrace{\langle\, \mathbf{h}_{n\mathbf{k}}\, |\, \mathbf{h}_{m(\mathbf{k}+\mathbf{b})}\, \rangle}_{=:\mathcal{M}_{nm}^{(\mathbf{k},\mathbf{b})}}, \tag{3.12}$$

where w_b is a scalar weight and $\mathbf{b}$ runs over the difference vectors between $\mathbf{k}$ and all neighboring $'vk$–points in $\mathcal{L}_{\mathbf{k}}$.

Eventually, it is possible to calculate a set of matrices $\mathcal{G}^{(\mathbf{k})}$ such that a unitary update

$$\mathcal{U}^{(\mathbf{k})} \to \mathcal{U}^{(\mathbf{k})} \cdot\ \exp\left(h\mathcal{G}^{(\mathbf{k})} \right) \tag{3.13}$$

results in a maximal increase of the objective for a fixed $0 < h \ll 1$. This matrix set is referred to as the gradient and the explicit form is

$$\begin{aligned} \mathcal{G}^{(\mathbf{k})} &= \frac{2}{|\mathcal{L}_{\mathbf{k}}|} \sum_{\mathbf{b}} w_b \left[\mathcal{R}^{(\mathbf{k},\mathbf{b})} - \mathcal{R}^{(\mathbf{k},\mathbf{b})} + \mathrm{i}\mathcal{T}^{(\mathbf{k},\mathbf{b})} + \mathrm{i}\mathcal{T}^{(\mathbf{k},\mathbf{b})} \right] && (3.14)\\ \text{with} \quad \mathcal{R}_{mn}^{(\mathbf{k},\mathbf{b})} &= \mathcal{M}_{mn}^{(\mathbf{k},\mathbf{b})} \mathcal{M}_{nn}^{(\mathbf{k},\mathbf{b})} && (3.15)\\ \text{and} \quad \mathcal{T}_{mn}^{(\mathbf{k},\mathbf{b})} &= \frac{\mathcal{M}_{mn}^{(\mathbf{k},\mathbf{b})}}{\mathcal{M}_{nn}^{(\mathbf{k},\mathbf{b})}} \left[\mathrm{Im} \ln \mathcal{M}_{nn}^{(\mathbf{k},\mathbf{b})} + \mathbf{b}\langle \mathbf{r} \rangle \right]. && (3.16) \end{aligned}$$

Derivations for this expression can be found in Refs. [29, 37].

The simplest way to find a local minimum of a function via its gradient is the proportional steepest–descent algorithm. In each iteration, a multiple of the gradient is

---

[3] The fact that $\mathbf{k}\cdot\mathbf{p}$–theory fails for 3D photonic crystals [42] already indicates that such systems might pose specific problems.

subtracted from the current state vector. The method is called "proportional" because the prefactor is constant throughout the minimization.

Applying this to the Wannier spread minimization problem, we find refined guesses for the final unitary mixing function by calculating

$$\mathcal{U}^{(\mathbf{k}),(i+1)} = \mathcal{U}^{(\mathbf{k}),(i)} \cdot \exp\left(-\alpha\ \mathcal{G}^{(\mathbf{k}),(i)}\right) \quad \text{while} \quad ||\mathcal{G}^{(i)}||^2 > h, \tag{3.17}$$

where i still enumerates the iterations, $\alpha > 0$ is the stability–creating deflation factor (typically $a \approx 0.5$), and h is the tolerance for deciding whether the approximation to the minimum is acceptable.

There is one single drawback of this method to minimize the spread of a set of Wannier functions. The term $\text{Im}\{\ln \mathcal{M}_{nn}^{(\mathbf{k},\mathbf{b})}\}$ (which was introduced by the approximation $\log x \approx x - 1$) is not defined if $\mathcal{M}_{nn}^{(\mathbf{k},\mathbf{b})} = 0$. This should not cause too much trouble because the Bloch modes of the same band at adjacent **k**–points should be as similar as possible to ensure good localization. Nevertheless, it can happen at a local minimum of the spread functional if the minimizer "trades" a discontinuity in one or two bands with very good localization in all other bands of a group. In such a case, the algorithm does not converge and cannot escape this situation.

3.2.2. Souza–Method for Artificially Isolating Band Groups

Well–localized (generalized) Wannier functions are the lattice–Fourier transform of generalized Bloch modes which are smooth and slowly varying functions with respect to **k**. At each **k**–point, the Bloch modes of a band group span a function space (a subspace of the full Hilbert space of physical states) and the problem of creating well–localized Wannier functions corresponds to finding a basis at each **k**–point such that the basis vectors vary smoothly and slowly with **k**. Obviously, this is not possible if the Hilbert spaces themselves do not vary smoothly and slowly. This is the case if the band group is entangled with other bands (cf. Fig. 3.1).

To overcome this problem and thereby allow for the construction of well–localized WFs for non–isolated band groups, the Marzari algorithm is not applied to (physical) Bloch modes but to generalized Bloch modes that were chosen in a preprocessing step such that the subspaces they span are as independent of **k** as possible. In our bottom–up approach (cf. Sec. 3.2.4), the band group only interacts with bands of higher frequency. We assume that there is a frequency window in which the mode information that is required to make the subspaces smooth with respect to **k** and the additional mode information from higher bands coexist. At each **k**–point, we now allow for a unitary mixing of all Bloch modes in this frequency range where the number of mixed Bloch modes may depend on **k**. The mixings are determined self–consistently such that the subspaces at neighboring **k**–points are as similar as possible. This subspace optimization method was developed by I. Souza [38] for electronic systems and adapted to photonic systems in our group [29] and the group of S. John in Toronto [41].

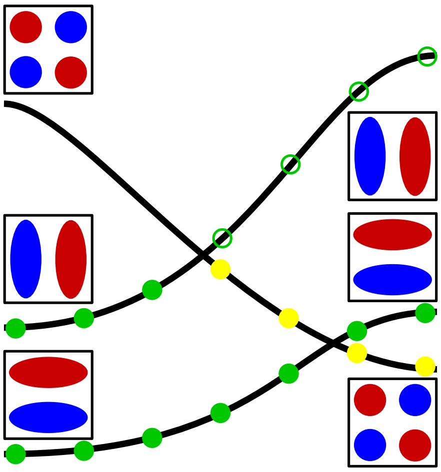

Figure 3.1.: *Sketch of subspaces formed by Bloch modes in a band structure. The solid green and yellow circles indicate the discretized bands of a band group that is entangled with a third band. Possible eigenmode patterns corresponding to the bands are added on the left and right hand side. The subspaces spanned by the first two bands (solid circles) at the discrete* **k***–points vary slowly except for the discretization points that surround the crossing of band two and three. The subspaces spanned by the modes indicated by green (solid and hollow) circles, however, vary slowly throughout the whole sketch.*

3.2.3. Trial Projection

The performance of both iterative algorithms tremendously depends on the initial Bloch mode mixing, which should already yield fairly well–localized and symmetric Wannier functions, even without further optimization. To this end, we project the Bloch modes into a set of trial functions (one trial per band, i. e. per desired Wannier function) and perform a symmetric orthonormalization as proposed in Refs. [37, 38].

This creates a set of generalized Bloch modes that form the best possible approximation to the trial functions if they are transformed via Eqn. (3.9). As a general rule of thumb, the spread minimization is more stable and requires less iterations the better the trial functions resemble the desired final Wannier functions, i. e. , the smaller the difference between the start configuration and the next local minimum is. Suitable trial functions can be as crude as randomly distributed Gaussians [37] or—a little more involved—Bloch modes that have been truncated around a high symmetry point [29]. Both approaches have the disadvantage that it is unclear which hybridization will come out or the subsequent minimization. Furthermore, this outcome can be very sensitive to slight modifications of the minimization parameters, which is very inconvenient.

For good, reproducible results the initial trial functions should be chosen such that they satisfy all formal constraints such as function center and symmetries (including

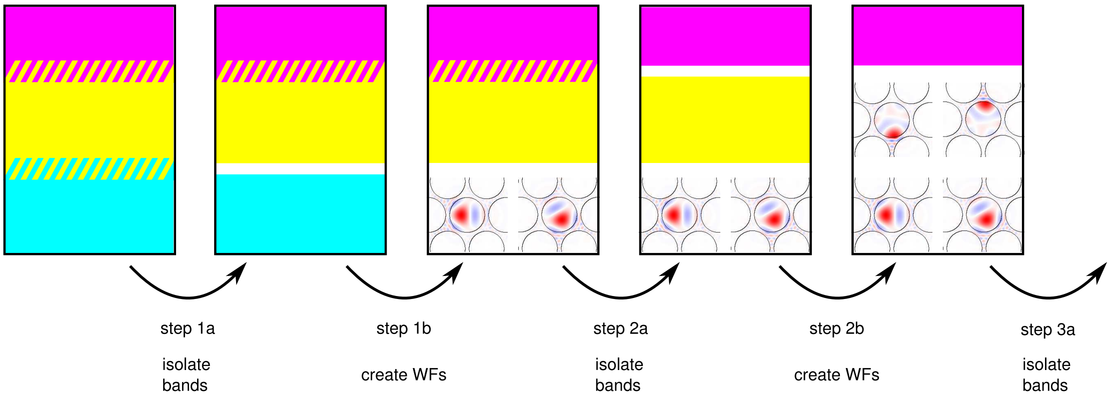

Figure 3.2.: *Sketch of the bottom–up approach for creating arbitrarily many photonic Wannier functions. Each color represents the mode information that is associated with a group of bands. The vertical direction of the boxes roughly corresponds to eigenfrequencies. In each stage, first, the lowest band group is isolated from the bands above using the Souza method. Then, a set of well–localized Wannier functions is created from this artificially isolated band group using the Marzari method and the process is repeated with the now lowest group of bands.*

the desired hybridization scheme). This requires a fair amount of theoretical analysis of the crystal's and the Bloch modes' symmetries and some work dedicated to the proper construction of trial functions [22].

3.2.4. Bottom–Up Approach

Sufficiently accurate Wannier–based calculations in the magnetic field formulation require a set of several dozen Wannier functions [22]. Unfortunately, it is not feasible to create all those functions together in the same spread minimization step because the number of degrees of freedom results in very long execution times, the minimizer's tendency to get stuck at a false minimum grows and, finally, because the resulting functions tend to be extremely hybridized. Although this favors excellent localization, it makes it practically impossible to reduce the basis to a subset of these functions as each function comprises contributions of Bloch modes from the full spectral range. Consequently, the large sets of basis functions are created as subsets that always correspond to an entangled group of (typically 15) bands and that is relatively weakly entangled with the bands above and below.

First, the lowest group (i. e. the band group with the lowest frequencies) is disentangled from the bands above using the Souza subspace optimization method. Then, the Marzari method is applied to this artificially isolated group, yielding the first subset of basis functions. Then the two substeps are repeated for the next lowest group of bands (cf. Fig. 3.2), resulting in the next basis subset. This bottom–up approach was developed by

Matthias Schillinger [29], it reduces the complexity of the minimization problems and allows to arbitrarily combine the subsets to form a reduced expansion basis.

The resulting Wannier functions strongly depend on which bands are grouped together. They must be chosen such that the groups can be efficiently disentangled from each other and that the symmetries of the constituent Bloch modes allow for hybridizations that are known to support tight localization [22]. Finally, there is a large contribution of guesswork and simple trial and error.

3.3. Advanced Minimization Strategies

While the Souza method for isolating the band groups from each other inherently relies on the fixed point iteration described in Sec. 3.2.2, the Marzari method poses a very general minimization problem for which a variety of strategies exist. The basic approach proposed in Ref. [37]—the steepest–descent method with a proportional step width—has already been presented in Sec. 3.2.1. However, it suffers from poor performance (cf. Sec. 3.3.1). To overcome this, I present two different methods, which exhibit much better performance at modest complexity: The gradient descent with adaptive step width Sec. 3.3.2 and the nonlinear conjugate gradient method Sec. 3.3.3 with a customized Wolfe–Powell line search Sec. 3.3.4. In Sec. 3.3.5, I finally discuss a randomized algorithm that was developed to find a very good minimum of the objective function even if no suitable trial functions are at hand. This is particularly useful for exploring completely new classes of photonic crystals.

3.3.1. Problems of the Basic Gradient Descent

Due to the singularities of the localization functional (which are introduced by approximating a linear function as the logarithm), the second derivative is not bounded. Thus, the prefactor must be chosen rather small to maintain stable operation in most cases. This has the disadvantage of causing very poor convergence rates once the well–behaved area around the minimum is reached. The performance of this method is shown in Tab. 3.1 and Fig. 3.3 in comparison to two other algorithms presented in the remainder of this section.

3.3.2. Adaptive Step Width Gradient Descent

The basic gradient descent method is suboptimal. The main problem is not that the update term is proportional to the gradient (this is a common choice) but that the proportionality factor α is fixed. We can expect much better performance, if it is adapted to the situations of each steepest–descent iteration, independently. Thus, we aim for a favorable sequence of step widths $\alpha^{(i)}$.

Without further knowledge about the optimization problem at hand, the standard approach (cf. Ref [43]) is to perform a line search along the direction defined by the gradient. This means that, within each iteration of the steepest–descent minimization, we

System	Band group	Method	Iterations	Updates	Gradients	Time
□	2–7	basic	7311	7311	7312	8.8 sec
□	2–7	adaptive	938	1876	939	2.4 sec
□	2–7	CG	70	126	197	0.3 sec
□	8–19	basic	16643	16643	16643	64 sec
□	8–19	adaptive	2270	4584	2271	16 sec
□	8–19	CG	184	313	498	1.4 sec
△	2–9	basic	57414	57414	57415	129 sec
△	2–9	adaptive	4129	8258	4130	16 sec
△	2–9	CG	222	487	710	1.6 sec
△	10–17	basic	77478	77478	77479	175 sec
△	10–17	adaptive	5458	10922	5459	21 sec
△	10–17	CG	198	422	621	1.4 sec
△	18–25	basic	26818	26818	26819	61 sec
△	18–25	adaptive	1908	3837	1909	7.6 sec
△	18–25	CG	126	259	386	1.0 sec
△	26–38	basic	498069	498069	498070	2961 sec
△	26–38	adaptive	41601	181862	41602	713 sec
△	26–38	CG	346	707	1054	5.0 sec

Table 3.1.: *Performance comparison between the three minimization methods discussed in this chapter for two types of photonic crystals. Both systems consist of circular air pores (relative radius* 0.45*) in silicon (*$\epsilon = 12$*) arranged in a square lattice (denoted by square symbols in the table) or a triangular lattice (triangles). In both cases, the calculations was performed on a reciprocal discretization of five* **k***–points per direction, the termination criterion was* $||\mathcal{G}||^2 < 10^{-12}$ *with respect to the same matrix norm in all three approaches. For the second system, the band grouping and window frequencies were taken from Ref. [29].*
The three methods are proportional steepest–descent method with deflation parameter (denoted with "basic"), steepest–descent method with Armijo line search (denoted with "adaptive") and Polak–Ribiére conjugate gradient method with customized strict Wolfe–Powell line search (denoted with "CG"). The free parameters of the methods were set to good standard values and not optimized to the specific requirements of the individual band groups. Nevertheless, in all cases, the adaptive method performs better than the basic approach and the conjugate gradients method is even faster. All performance data were obtained from my own implementation of the Wannier generation algorithms run on a single core of my workstation (Intel Core2 Q9400 @ 2.66GHz with 8GB of memory). The execution times cover the full run time of the minimizer including loading the precomputed overlap matrices and storing the final mixing function.

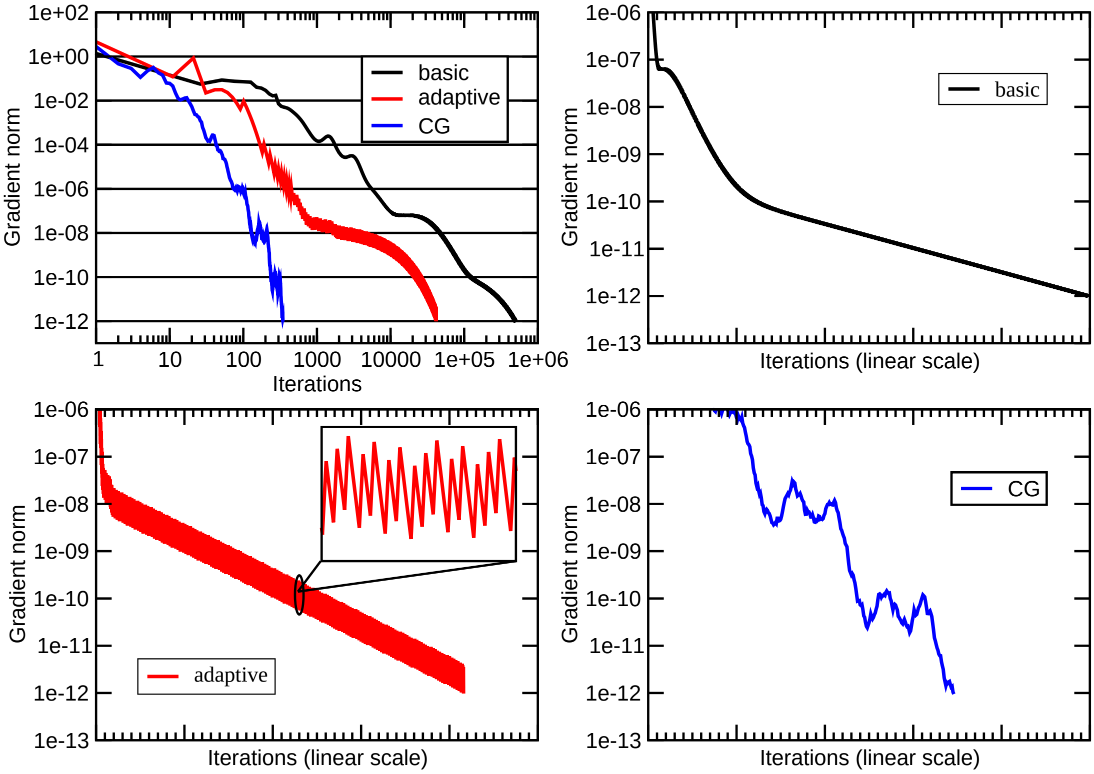

Figure 3.3.: *Convergence behavior of the three minimization methods discussed in this chapter. In this plot, the band group* 26–38 *in the triangular lattice of air pores in silicon (cf. last three lines in Tab. 3.1) was minimized using the proportional steepest–descent method (black curves), the steepest–descent method with Armijo rule (red curves) and the Polak–Ribiére method with customized strict Wolfe–Powell line search (blue curves). It is clearly visible, that the CG method converges much (up to three orders of magnitude) faster than the other ones, although the asymptotic convergence of the gradient descent methods is exponential. Most likely, the local objective minimum forms a narrow valley which is known to be a serious challenge for steepest–descent methods. Oscillations of the gradient norm as visible in the red curves are a typical symptom of this problem.*

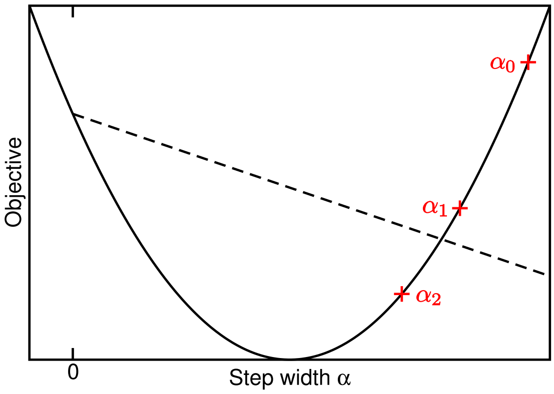

Figure 3.4.: *Sketch of the Armijo method. The objective function restricted to the search line is depicted as a solid black line, the acceptance line as a dotted one. The dotted line's slope is a fraction of the slope of the tangent (not shown) at the start point of the line search (denoted with "0"). On the right hand side, there are the initial guess* α_0 *for the step width and two subsequent refinements with a reductions factor of* $\beta = 0.85$*.* α_0 *yields a worse result than the start point and, thus, is not an acceptable step width.* α_1 *does improve the objective but stays above the acceptance line.* α_2*, finally, lies below the acceptance line and is an acceptable step width.*

perform an iterative minimization procedure to find the optimal (or at least a sufficiently good) value of $\alpha^{(i)}$ for this particular gradient. Thus, we consider the objective function restricted to the line along the gradient and aim for the minimum of this restriction.

This optimal step width can be analytically determined for very special cases such as a quadratic objective function. In general, however, one settles for a step width that is guaranteed to reduce the objective and often is close to the optimal step width. Typical general–purpose methods for this are the Armijo method or the Wolfe–Powell method [43]. In practice, the Armijo method proved to considerably increase the convergence rate of the steepest–descent algorithm.

This approach is illustrated in Fig. 3.4. The basic idea is to start with an initial guess α_0 for the step width and to check whether it (sufficiently) improves the objective. If it does not, the guess is reduced by a fixed factor β (which is a parameter of the method):

$$\alpha^{(i),(j)} = \alpha_0 \beta^j, \tag{3.18}$$

where $j = 0, 1, 2, \ldots$ enumerates the line–search iterations and $\beta = 0.7$ is a good choice. These two steps are repeated until an acceptable step width is found. This method terminates as long as the directional derivative of the objective in search direction is negative (always true for the steepest–descent).

One might expect that every step width that reduces the objective function is acceptable. However, in this case it would be possible that an arbitrarily large step width comes with only an arbitrarily small improvement, which is not sufficient to guarantee any sort of convergence rate. Thus, we define that a step is acceptable only if the ratio between objective improvement and step width is larger than a fraction of the directional derivative:

$$\Omega(\alpha^{(i),(j)}) < \Omega(0) - \alpha^{(i),(j)} \sigma ||\mathcal{G}||, \tag{3.19}$$

Here, $-\sigma||\mathcal{G}||$ is the slope of a line through $\Omega(0)$, which decides whether a possible step width is acceptable. Ref. [43] recommends to restrict σ to the interval $(0, 0.5)$ because the minimum of a quadratic objective (which is a good approximation close to the minimum) should not be excluded, as this would slow down the process unnecessarily.

Introducing this acceptance line has an additional advantage if the objective restricted to the search line has more than one (local) minimum. In this case, the minimum corresponding to the smallest step width (the minimum adjacent to the starting point of the line search) is preferred because the other ones lie above the acceptance line unless they are much better than the former one or the acceptance line is rather flat. This suggests to extend the allowed parameter range to $0 < \sigma < 1$ if the minimization is known to tend to run into undesired local minima.

In the Armijo method, the step width has an upper bound defined by α_0, which can be independently chosen for each line search run, potentially based on some clever heuristic (provided one has the patience to develop one). The choice of the initial guess is not extremely critical, though. If it is too small (between the start point and the adjacent minimum), more gradient descent iterations are required. If it is too large, many refinements and checks are required. In either case the convergence rate is reduced but no severe problems occur. For simplicity, I use the directional derivative's norm multiplied by some factor (typically 30). Effectively, this behaves like the basic steepest–descent method with a very large step width prefactor that is automatically reduced whenever problems would occur. This combination of steepest–descent and adaptive step width performs much better than the simple method as depicted in Tab. 3.1. However, it is well known that optimization methods that perform line searches only along the gradient direction have difficulties with finding the minimum of a narrow valley. In general, this is the case during the generation of maximally localized Wannier functions and the symptoms are nicely visible in Fig. 3.3.

3.3.3. Conjugate Gradients

Let us forget about Wannier functions and Bloch modes for a moment. As already mentioned, the performance problems of steepest–descent methods close to narrow valleys are well known and the most popular approach to overcome them is the conjugate gradient (CG) method. In general, one considers the task of minimizing an objective function Ω

$$\Omega(\vec{x}) : \mathbb{R}^n \mapsto \mathbb{R} \tag{3.20}$$

over the Euclidean space of dimensionality n. Said convergence problem can be traced down to a zig–zag motion of the minimizer, i. e. , the steepest–descent method repeatedly provides similar descent directions. This can be avoided by projecting the gradient onto a subspace that does not include any of the previous directions. Naively, one could require subsequent descent directions to be orthogonal, which unfortunately does not perform well, either. Instead, the minimum often is approximately quadratic and, therefore, the Hessian $\mathcal{J}$ is approximately constant. One chooses the conjugate space with respect to $\mathcal{J}$:

$$\vec{d}^{\,j} \cdot \left(\mathcal{J} \cdot \vec{d}^{\,i} \right) = 0 \quad \forall_{i<j}, \tag{3.21}$$

where $\vec{d}^{\,i}$ is the descent direction of the i–th iterations step.

If the objective is exactly quadratic, i. e. ,

$$\Omega(\vec{x}) = \frac{1}{2}\vec{x} \cdot (\mathcal{A}\vec{x}) - \vec{b} \cdot \vec{x}, \tag{3.22}$$

the Hessian is trivially known ($\mathcal{J} = \mathcal{A}$) and Eqn. (3.21) is fulfilled by a recursively defined sequence of directions

$$\vec{d}^{(i)} = -\vec{\nabla}\Omega(x^{(i)}) + \frac{|\vec{\nabla}\Omega(x^{(i)})|^2}{|\vec{\nabla}\Omega(x^{(i-1)})|^2}\vec{d}^{(i-1)}, \tag{3.23}$$

with $\vec{d}^{(0)} = \nabla\Omega(x^{(0)})$. With an appropriate choice of the step width, this is the linear conjugate gradient method, which is often used to solve systems of equations $\mathcal{A}\vec{x} = \vec{b}$ if $\mathcal{A}$ is large symmetric and/or not explicitly known.

This still works for a not strictly quadratic minimum by simply inserting the nonlinear gradient. However, it is well known that this so–called Fletcher–Reeves method [43] is not the most performant way to adapt the conjugate gradients method to nonlinear problems. Another way, which very often converges faster, is the Polak–Ribière [43] method based on the descent direction sequence

$$\vec{d}^{(i)} = -\vec{\nabla}\Omega(x^{(i)}) + \frac{\left(\vec{\nabla}\Omega(x^{(i)}) - \vec{\nabla}\Omega(x^{(i-1)})\right) \cdot \vec{\nabla}\Omega(x^{(i)})}{|\vec{\nabla}\Omega(x^{(i-1)})|^2}\vec{d}^{(i-1)} \tag{3.24}$$

to approximate the conjugate gradients.

In contrast to the linear conjugate gradients method for solving systems of linear equations[4], nonlinear conjugate gradient methods only provide a descent direction and require a subsequent line search for determining an adequate step width. The Armijo method presented in the previous section is one possibility. However, in the linear conjugate gradients method, the recursive formula Eqn. (3.23) only provides conjugate descent directions if the step width is optimally chosen. One would expect that this condition remains approximately valid for nonlinear conjugate gradients. The Armijo method provides only a rough approximation to the minimum of the line search problem. The customized strict Wolfe–Powell method (cf. Sec. 3.3.4) typically provides a much better approximation and yields much faster convergence[5].

Unfortunately, we cannot directly apply the Polak–Ribiére method to our problem at hand because it requires all possible configurations to form a vector space with an associated scalar product. Clearly, the unitary mixing functions $\mathcal{U}^{(\mathbf{k})}$ do not. Furthermore, a unitary matrix is redundant in the sense that roughly half of the entries are defined by the other half and the unitarity property, so we cannot simply map the matrix on a state vector which consists of all matrix elements.

However, we can express any unitary matrix in terms of an anti–Hermitian one [37]:

$$\mathcal{U} = \exp\left(\mathcal{W}\right), \quad \mathcal{W}^\dagger = -\mathcal{W}. \tag{3.25}$$

The reader might argue that this substitution is not unique and, thus, the parameter space is overdefined. This is true but does not pose any practical problems: The objective is continuously replicated across the configuration space. More specifically, it becomes periodic with respect to the eigenvalues of $\mathcal{W}$. Furthermore, the gradients for equivalent anti–Hermitian matrices are identical[6].

---

[4] A quadratic objective function is one of the rare examples where the optimal step width can be determined exactly.

[5] Curiously, this line search method does not perform very well with the steepest–descent direction method. Probably this is because the accuracy of the line search method does not influence the quality of the descent directions and, thus, the increased numerical effort to find the optimal step width is wasted.

[6] This might suggest that Ω is even periodic with respect to $\mathcal{W}$ itself, but I did not put any effort in this question as it is irrelevant for the WF generation.

Now, we can regard the real and imaginary parts of all matrix elements in the upper right half (excluding the diagonal) and the imaginary parts of the diagonal entries as independent real–valued degrees of freedom. Thereby, we can identify each anti–Hermitian matrix with exactly one member of the Euclidean space $\mathbb{R}^{N^2}$, where N is the number of rows in $\mathcal{W}$. Consequently, we can identify the vector addition with the standard matrix addition[7] and the scaling of a vector with the multiplication of the same scalar with the respective matrix. Finally, I define a compatible inner product as:

$$\langle \mathcal{W} | \mathcal{V} \rangle := \sum_{n,m \geq n} \mathrm{Re}\Big\{ (\mathcal{W}_{nm})^* \mathcal{V}_{nm} \Big\}. \tag{3.26}$$

It is noteworthy, that this inner product is not sesquilinear although the matrices involved are complex and that the induced norm is none of the usual matrix norms. Using this, we can formulate the Polak–Ribiére method for our problem:

$$\begin{aligned} \mathcal{D}^{(\mathbf{k}),(0)} &= -\mathcal{G}^{(\mathbf{k}),(0)}, & (3.27)\\ \mathcal{D}^{(\mathbf{k}),(i)} &= -\mathcal{G}^{(\mathbf{k}),(i-1)} + \frac{\langle \mathcal{G}^{(\mathbf{k}),(i)} - \mathcal{G}^{(\mathbf{k}),(i-1)} | \mathcal{G}^{(\mathbf{k}),(i)} \rangle}{\langle \mathcal{G}^{(\mathbf{k}),(i)} | \mathcal{G}^{(\mathbf{k}),(i)} \rangle} \cdot \mathcal{D}^{(\mathbf{k}),(i-1)}, & (3.28)\\ \mathcal{U}^{(\mathbf{k}),(i+1)} &= \mathcal{U}^{(\mathbf{k}),(i)} \exp\Big(\alpha^{(i)} \mathcal{D}^{(\mathbf{k}),(i)} \Big). & (3.29) \end{aligned}$$

Usually, there would be no point in implementing this by hand and one would use a readily available[8] optimization library. However, performing the mapping between the $\mathcal{W}$–matrices and the real–valued vector triples the memory consumption (in Eqn. (3.29), there is no need to store the $\mathcal{W}$–matrices) and considerably increases the execution time as compared to a custom implementation of Eqn. (3.29). Therefore, it is indeed a good idea to program this by hand.

Conjugate gradients provide an efficient minimizer whose implementation is not too complex. The next logical step in the quest for even faster optimization would be the class of Newton and Quasi–Newton methods. As the name suggests, they require the second derivative (Hessian) of the objective or (in the latter case) at least an approximation extracted from subsequent gradient evaluations. However, our problems typically have several thousand degrees of freedom, which corresponds to a huge Hessian and for sure excludes true Newton methods. In fact, there are some Quasi–Newton methods that do not require the approximation of the Hessian to be stored explicitly. However, the number of iterations and the execution time with the CG method is already rather low (cf. Tab. 3.1 and Fig. 3.1), so I doubt that the time required for programming and testing would pay off.

3.3.4. Customized Strict Wolfe–Powell Method

For conjugate gradient minimization, the Armijo line search method turned out to be suboptimal. Most likely, this is because this method is somewhat crude and tends to

---

[7] This is good news. Otherwise, the expression for the objective's gradient above would be useless.

[8] . . . and well–tested—it's amazing how many mistakes can be incorporated in the implementation of this simple formula. . .

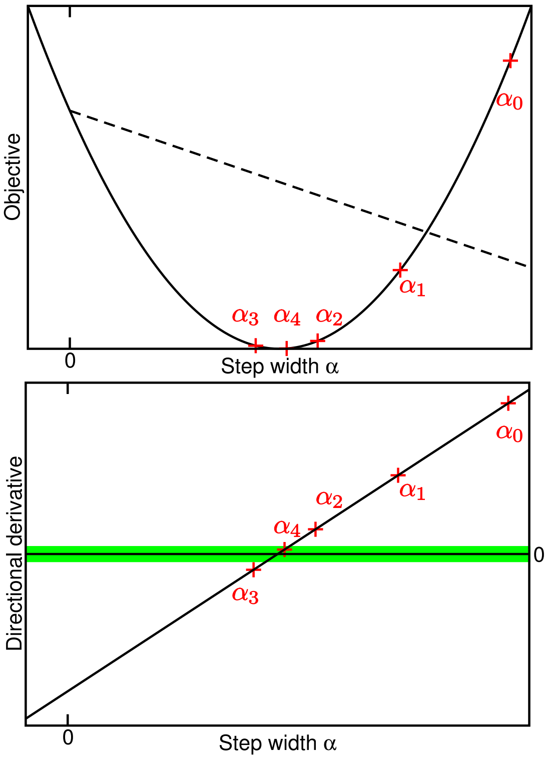

Figure 3.5.: *Sketch of the strict Wolfe–Powell line search. The objective function restricted to the search line and its derivative with respect to the step width* α *are depicted as solid lines in the top and in the bottom panel, respectively. The dotted line is an acceptance line in analogy to the Armijo line search method. The green area indicates the range of acceptable derivatives. This sketch illustrates four subsequent iteration steps, denoted* α_1 *through* α_4 *with* α_0 *being the start point. The first three steps form an exponential decrease until a sign change of the directional derivative is detected (i. e. until the minimum has been passed) below the Armijo line. Once a minimum is roughly located, a bisection series of step widths is formed to find a step width that is sufficiently close to the minimum. In this sketch, only one bisection step is required as* α_4 *is already acceptable.*

choose too large step widths. However, recursion Eqn. (3.23) only provides exactly conjugate directions if the optimal step width is chosen in each iteration. Therefore, we need a line search method that typically finds a much better approximation to a local minimum of the objective restricted to the search line.

The strict Wolfe–Powell method (defined in section 5.2 of Ref. [43]) aims for a minimum of $\Omega(\alpha)$ rather than just some step width that reduces the objective. An acceptable approximation to the minimum is a step that:

1. Sufficiently reduces the objective. This is ensured by an acceptance line criterion already known from the Armijo line search (Eqn. (3.19)). From now on, I will refer to this acceptance line as the "Armijo–line".

2. Sufficiently reduces the directional derivative $\partial_\alpha \Omega$ along the search line. An acceptable step width must fulfill the criterion

   $$\left|\partial_\alpha \Omega(\alpha)\right| < \rho \left|\partial_\alpha \Omega(0)\right| , \tag{3.30}$$

   where ρ is a small positive number.

This is illustrated in Fig. 3.5.

My implementation for finding such a step width is a customization of the sample implementation "algorithm 6.5" of Ref. [43]. The line search is divided into two stages:

1. Roughly locate the minimum,

2. Track down the minimum.

Stage One — Locating

The first stage is very similar to the Armijo method: Let there be an initial guess α_0. If $\Omega(\alpha_0)$ is above the Armijo–line or if $\partial_\alpha \Omega(\alpha_0) > 0$, the step width is reduced by multiplying it with a factor β until the directional derivative's sign changes from positive for α_{n-1} to negative for α_n and $\Omega(\alpha_{n-1})$ and $\Omega(\alpha_n)$ both are below the Armijo–line.

If $\Omega(\alpha_0)$ is below the Armijo–line and if $\partial_\alpha \Omega(\alpha_0) < 0$, the step width is increased by dividing it by the factor β until the directional derivative's sign changes from negative for α_{n-1} to positive for α_n and $\Omega(\alpha_{n-1})$ and $\Omega(\alpha_n)$ both are below the Armijo–line.

If this is not possible for some reason, the whole minimization is aborted. In particular, the second case may fail occasionally. Otherwise, there is at least one acceptable minimum in the interval (α_{n-1}, α_n).

Stage Two — Tracking

Now, a local minimum of $\Omega(\alpha)$ is roughly localized and can be tracked down via bisection. We define a sequence of intervals (a_i, b_i), which contain the minimum. The bisection is initialized by

$$a_0 = \min\{\alpha_{n-1}, \alpha_n\} \quad \text{and} \quad b_0 = \max\{\alpha_{n-1}, \alpha_n\}. \tag{3.31}$$

In each bisection iteration, the interval boundaries are refined according to

$$a_{i+1} = \quad a_i \text{ if } \partial_\alpha\Omega(c_i) > 0 \quad \text{or} \quad c_i \text{ otherwise,} \tag{3.32}$$

$$b_{i+1} = \quad c_i \text{ if } \partial_\alpha\Omega(c_i) > 0 \quad \text{or} \quad b_i \text{ otherwise,} \tag{3.33}$$

where the interval center $c_i = \frac{1}{2}(a_i + b_i)$ becomes one of the new interval boundaries. In order to avoid erratic behavior, the objective is evaluated as well and compared to the Armijo–line.

This is a very simple procedure which slowly converges to the minimum. However, evaluating the objective and its gradient are rather expensive tasks, which dominate the computational cost of the line search. Thus, even more sophisticated choices for c_i rather than the simple interval center will not reduce the performance of the line search as long as they do not require any additional evaluations of Ω or $\partial_\alpha\Omega$. On the other hand, this can drastically speed up the convergence if the c_i proposed by the heuristic in most cases lies close to the minimum.

To this end, I approximate $\Omega(\alpha)$ as a cubic polynomial

$$P(\alpha) = A\alpha^3 + B\alpha^2 + C\alpha + D \tag{3.34}$$

based on the previous four[9] evaluations of Ω. Its derivative is quadratic, thus the extrema of the interpolation can be easily calculated:

$$c_i = \frac{1}{6A}\left(-2B \pm \sqrt{4B^2 - 12AC}\right). \tag{3.35}$$

If one of them is a minimum, which lies inside the interval (a_i, b_i), it is chosen as the new boundary candidate. If the cubic polynomial does not have a minimum or if the minimum lies outside (a_i, b_i), the simple bisection rule is used as a fall–back.

---

[9] A similar approximation based on the *two* previous evaluations of Ω and of $\partial_\alpha\Omega(\alpha)$ is possible as well.

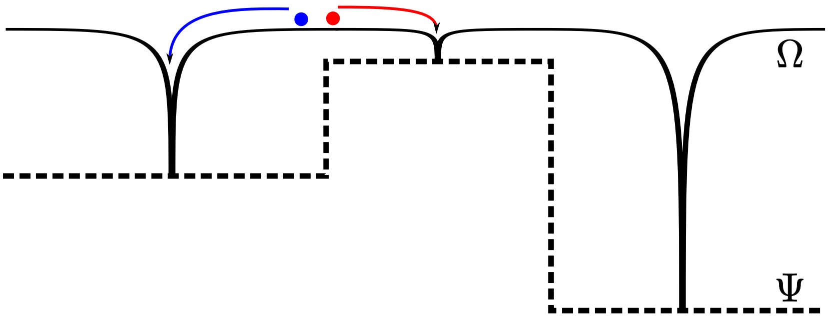

Figure 3.6.: *Sketch of how the reduced objective for the simulated annealing is defined. The reduced objective Ψ for a certain configuration is defined as the final value of Ω in a complete CG minimization starting from that very point. Ψ is discontinuous because neighboring points (red and blue dot) in the configuration space can only have different values if the CG minimizer approaches different local minima of Ω.*

3.3.5. Add–On: Hybrid Simulated–Annealing / Gradient Minimization

Proper trial functions are crucial for nice, well–hybridized and well–localized Wannier functions. However, constructing them requires knowledge about what types of Wannier functions can be obtained. The rigorous approach is a complete symmetry analysis [22]. Unfortunately, this is already a tedious and time–consuming job for two–dimensional symmorphic space groups. For three–dimensional non–symmorphic groups (e. g. the woodpile structure), we expect even more complications and have not yet performed such an analysis.

As a last point in my discussion of spread minimization strategies, I present a randomized algorithm which does not find the bottom of a local minimum defined by sophisticated trial functions. Instead, it aims to find the global minimum of the objective even if the start point is very poor.

To this end, I employ the simulated annealing (SA) method [44], which is a specific type of randomized optimization algorithms. It is physically motivated by the fact that slowly cooling down a thermodynamic system (such as a work piece made of metal) results in a final state with nearly minimal free energy. In the simulated annealing optimization method, a configuration is repeatedly improved by perturbing it and checking whether the perturbed configuration yields a smaller objective $\Omega^{\text{perturbed}}$ than the unperturbed Ω^{orig} one. If this is the case, the perturbed configuration is accepted. If not, the perturbed configuration is accepted with a probability which decreases exponentially with the difference in the objective function:

$$\mathrm{P} = \exp\left(-\beta(\Omega^{\text{perturbed}} - \Omega^{\text{orig}})\right) . \tag{3.36}$$

This resembles the exponentially suppressed occupation of energetically unfavorable states in a thermodynamic system. Accordingly, the suppression exponent β is identified with an inverse temperature and slowly increased (the system is cooled down). This exponential factor allows the minimizer to pass barriers and, thus, to find a good minimum. As the system is cooled down, the minimizer is caught in the vicinity of this minimum.

Our objective Eqn. (3.10) is not well–suited for a minimization via simulated annealing because the vast majority of possible configurations (i. e. sets of unitary matrices) yields extremely poorly localized WFs. It merely resembles a flat surface with narrow, deep pits. Consequently, the minimizer has trouble to identify the vicinity of a good minimum if the temperature is high and cannot overcome the barriers once the temperature is low.

However, the simulated annealing method and the CG method can be efficiently combined. In my approach, each simulated annealing step comprises a complete CG minimization. Starting from an initial configuration ($\mathcal{U}^{\mathrm{A}}$), the spread is minimized using the CG method, resulting in a spread Ω^{A}. Then, the initial configuration is perturbed (denoted $\mathcal{U}^{\mathrm{B}}$) and another CG minimization is performed starting from this perturbed configuration. If $\Omega^{\mathrm{B}} \leq \Omega^{\mathrm{A}}$, the perturbed configuration is accepted as the new current configuration. Otherwise, the perturbed configuration is accepted with the typical exponentially suppressed probability

$$\mathrm{P}_{\mathcal{U}^{\mathrm{A}} \to \mathcal{U}^{\mathrm{B}}} = \exp\Big(-\beta(\Omega^{\mathrm{B}} - \Omega^{\mathrm{A}})\Big) \,. \tag{3.37}$$

This corresponds to applying the SA algorithm to a skyline–shaped reduced objective Ψ as illustrated in Fig. 3.6. The advantage of this reduced objective is that the SA algorithm can freely move through the catchment area of a local minimum. This allows for much greater perturbations and higher temperatures as compared to applying the SA method to the original objective because the SA algorithm does not have to actually find the minimum. It is merely used to find the vicinity of the best available minimum. This approach has proven useful in situations where it is unclear which type of Wannier functions can be created and, hence, no suitable trial functions are available. In fact, it even works (in many cases) if no trial projection was performed at all. The severe disadvantage is that there is only little possibility to influence the outcome because of its random nature.

3.3.6. Secondary Objectives: Imaginary Part Reduction

Up to this point, we discussed the basic approach for constructing maximally localized photonic Wannier functions including some optimizations of the minimization procedure. However, small spread is not the only figure of merit for Wannier functions that are intended as a basis for numerical simulations.

First, for most practical physical systems, it is possible to construct purely real–valued Wannier functions. This is a worthwhile property as it reduces the memory required for storing the basis by a factor of two and leads to real–symmetric system matrices (instead of complex–Hermitian ones) in the WF–based simulation if the system is lossless. Again, this reduces the memory consumption of the calculation by a factor of two but, more importantly, it allows us to use more specialized matrix algorithms (e. g. eigenvalue solvers). This can considerably reduce the execution time of the simulation.

Second, the coupling between basis functions at different lattice sites depends on the localization properties. Even very tightly localized (in terms of the spread functions) basis functions can suffer from great long–range interaction if only one of them contains a slowly decaying tail. This results in unnecessarily dense system matrices and higher resource consumption of the simulations.

Whether Wannier functions are real–valued or exponentially localized is related to certain symmetries (see Ref. [22] for the connection to exponential localization). In this section, I present one way to incorporate symmetry constraints in the basic localization minimizer. In this section, I apply this ansatz to the problem of minimizing the imaginary parts of Wannier functions; I briefly present some groundwork for applying it to point symmetries in the next section.

The following discussion about an auxiliary penalty term for reducing the imaginary parts of WFs is a modified version of our manuscript Ref. [45].

Invariance under time–reversal guarantees that counter–propagating Bloch modes may be chosen to fulfill

$$\mathbf{H}_{n\mathbf{k}}(\mathbf{r}) = \mathbf{H}^{*}_{n(-\mathbf{k})}(\mathbf{r}). \tag{3.38}$$

If this holds, the Wannier functions according to Eqn. (3.1) are real–valued, which is not affected by a unitary mixing (cf. Eqn. (3.3)) that fulfills the condition

$$\mathcal{U}^{(\mathbf{k})}_{nm} = \left(\mathcal{U}^{(-\mathbf{k})}_{nm}\right)^{*}. \tag{3.39}$$

This relation could be used to eliminate half of the degrees of freedom in the minimization procedure. For instance, the minimization could be performed only on half of the Brillouin zone and the mixing matrices on the other half could be derived afterward by complex conjugation. This approach reduces the dimensionality of the minimization problem but requires proper boundary conditions at the border of the reduced **k**–space domain. While this is fairly simple in the case of time–reversal symmetry, it can become rather involved for point symmetry groups.

A second possibility is to start the minimization with a set of Bloch modes which already yields WFs with the desired symmetry (but poor localization) and to strictly symmetrize the gradient in each iteration before applying it. In the specific case of time–reversal, the trial projection can be performed on half of the Brillouin zone and the other half of the mixing function can be obtained by complex conjugation, which ensures a purely real–valued starting point. To maintain the "real–valuedness" during the minimization, the gradient can be calculated throughout the whole Brillouin zone and symmetrized according to

$$\mathcal{G}^{(\mathbf{k})} \to \mathcal{G}^{(\mathbf{k})} - \mathcal{G}^{(-\mathbf{k})}. \tag{3.40}$$

A third approach is to weakly enforce the desired symmetry properties by adding an auxiliary penalty term to the minimization objective. The penalty term measures the deviation from the desired symmetry and can be weighted with respect to the localization functional. This ansatz mimics strict gradient symmetrization if the weight is sufficiently large, but provides further control over the minimization procedure. In particular, it is possible to ramp up the symmetry constraint once the minimizer has reached the vicinity of the desired spread minimum. However, this comes at the expense of requiring a (preferably closed) expression for the gradient of the auxiliary penalty term.

Besides the availability of an analytical gradient, the penalty term must be a real–valued, scalar function that is independent of $\mathbf{r}$, $\mathbf{k}$ and n. I choose the accumulated squared norms of the imaginary parts of all Wannier functions

$$\Omega_{\mathrm{im}} = \sum_n \int \mathrm{d}r \, \Big[\mathrm{Im}\mathbf{W}_{n\mathbf{0}}(\mathbf{r}) \Big]^2 \tag{3.41}$$

$$= \sum_n \frac{1}{2} \Big[1 - \mathrm{Re}\langle\, {\mathbf{W}_{n\mathbf{0}}}^* \,|\, \mathbf{W}_{n\mathbf{0}} \,\rangle \Big], \tag{3.42}$$

where the integration is carried out over the full real space. The corresponding gradient reads:

$$\mathcal{G}^{\mathbf{k}}_{\mathrm{im}} = -\frac{1}{2} \Big[(\mathcal{U}^{-\mathbf{k}})^T \cdot \mathcal{U}^{\mathbf{k}} - (\mathcal{U}^{\mathbf{k}})^\dagger \cdot (\mathcal{U}^{-\mathbf{k}})^* \Big]. \tag{3.43}$$

The main steps for calculating this expression can be found in Appendix A.1. Finally, we add the weighted penalty term and the basic localization functional and obtain the total minimization objective and its derivative

$$\Omega_{\mathrm{tot}} = \Omega + \lambda \Omega_{\mathrm{im}}, \tag{3.44}$$

$$\mathcal{G}^{\mathbf{k}}_{\mathrm{tot}} = \mathcal{G}^{\mathbf{k}} + \lambda \mathcal{G}^{\mathbf{k}}_{\mathrm{im}}, \tag{3.45}$$

respectively, where λ is the real–valued penalty weight.

To illustrate the performance of this approach I apply it to the first two bands of an fcc–woodpile structure consisting of silicon ($\epsilon_r = 12$) rods with square cross section in air. As I will explain in Sec. 3.4, it is very challenging to construct good Wannier functions from these bands of 3D photonic crystals. This is one of the relatively rare occasions where Wannier functions tend to acquire really large imaginary parts and thus it is perfectly suited for illustrating the method. However, reducing the imaginary part is equally useful in cases where its “natural” value is around 10^{-5}.

As the first two bands are separated from the rest of the band structure (cf. Fig. 3.11) by a band gap, there is no need to perform a Souza disentanglement step. For initializing the Marzari minimizer, I projected the Bloch modes onto Gaussians centered at Γ. As the trial functions were real–valued, the imaginary–part at the starting point vanish. In Fig. 3.7, I plotted the sum over the combined norms of the imaginary parts of both bands (Ω_{im} from Eqn. (3.41)) with (red and blue curves) and without (black line) the penalty term. Using only the basic algorithm (i.e. $\lambda = 0$), the WFs quickly acquire and maintain considerable imaginary parts.

Setting the penalty weight λ to a fixed value (e.g. $\lambda = 200$, corresponding to the red curve in Fig. 3.7) improves the situation as the imaginary parts are drastically reduced

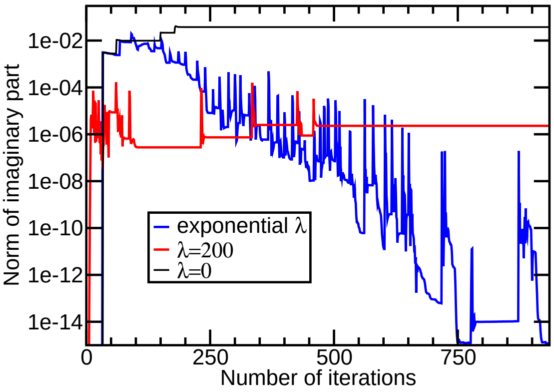

Figure 3.7.: *Performance of the auxiliary penalty term applied to the first two bands of a woodpile structure. The black curve depicts the unweighted total imaginary part norm* Ω_{im} *if the minimization is performed with the basic localization functional* Ω *only. The red and blue lines depict the same quantity for constant* $\lambda = 200$ *or exponentially growing* λ*, respectively. The minimization was performed using the conjugate gradient method.*

(in this example to values around 10^{-6}). This value of λ was chosen to illustrate that the auxiliary penalty term causes a finite reduction of the imaginary part; it is not extremely good in the sense that the residual distortion is still noticeable. However, finding the optimal value for λ is not straightforward. On the one hand, a large penalty weight guarantees that the imaginary part becomes negligibly small. On the other hand, a large λ is problematic in the beginning of the minimization procedure when the step width typically is large.

It is best to start with a small penalty weight and increase it exponentially. This way, the minimizer is nearly free in the beginning and can settle to a good local minimum. Once it cannot easily escape from this minimum, the configuration space is gradually restricted until the desired degree of symmetry is reached. As an example, I started with $\lambda = 0.1$ and increased it by 2% per iteration. Larger values than $\lambda = 10^5$, should be avoided because round–off errors start to become problematic and it is advisable to stop the exponential growth when this boundary is exceeded. In this example, this happened after about 700 iteration steps. The blue curve in Fig. 3.7 illustrates that this method efficiently minimizes the imaginary parts. The (real parts of the) resulting Wannier functions are depicted in Fig. 3.8.

Irrespective of whether a fixed or a growing penalty weight is used, the parameters cannot be determined a priori. Instead, they have to be empirically determined in several minimization runs with different parameters, where a rough estimate can be made based on experience.

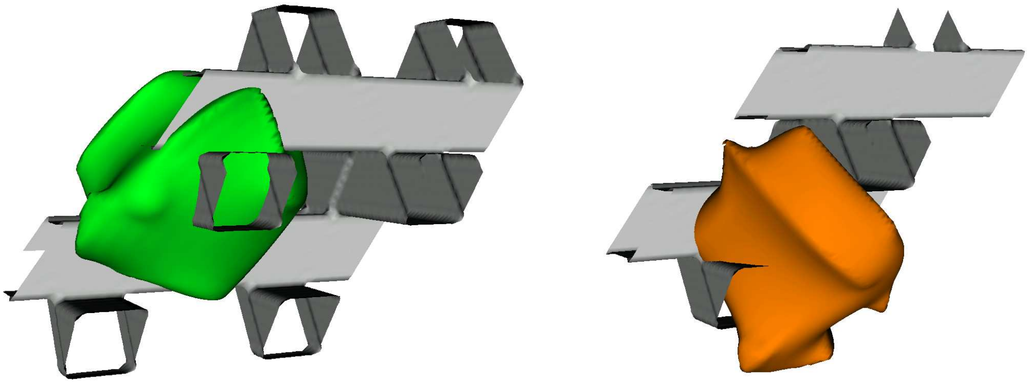

Figure 3.8.: *Band one and two of the final Wannier functions (green and orange) generated for the woodpile structure (outline depicted in gray). Depicted are 50% isocontours of the magnetic field with respect to the Wannier functions' field maxima. For these functions, I employed the auxiliary penalty term Ω_{im} with exponentially growing weight λ, corresponding to the blue curve in Fig. 3.7.*

3.3.7. Outlook: Point Symmetries

As already alluded to, the basic idea of the previous section can be applied to point symmetries as well. In this section, I would like to propose how this could be accomplished.

Let us assume that the target Wannier functions form a representation [22] of the photonic crystal's point group $\mathcal{P}$. This means that for each element $\hat{T}$ of the point group, there exists a matrix (the representation matrix) $\mathcal{D}^{(\hat{T},\mathrm{WF})}$ which can be chosen to be unitary and reflects the action of the symmetry operation:

$$\hat{T}|\,\mathbf{W}_n\,\rangle = \sum_m |\,\mathbf{W}_m\,\rangle \mathcal{D}^{(\hat{T},\mathrm{WF})}_{mn}, \tag{3.46}$$

where only Wannier functions at the lattice origin $\mathbf{R} = \mathbf{0}$ are taken into account.

On the other hand, $\hat{T}$ also relates the Bloch modes at $\mathbf{k}$ and $\hat{T}\mathbf{k}$ via another (unitary and $\mathbf{k}$–point dependent) point group representation[10] $\mathcal{D}^{(\hat{T},\mathrm{BF},\mathbf{k})}$:

$$\hat{T}|\,H_{n\mathbf{k}}\,\rangle = \sum_m |\,H_{m\hat{T}\mathbf{k}}\,\rangle \mathcal{D}^{(\hat{T},\mathrm{BF},\mathbf{k})}_{mn}. \tag{3.47}$$

If $\mathbf{k}$ is a general point, the matrices can be chosen to be the unit matrices. If, however, $\mathbf{k}$ is mapped onto itself by some symmetry operations, the representation matrices may become nontrivial.

---

[10] For the sake simplicity, I only consider scalar–valued Bloch modes here. This is also justified because plain 2D problems can always be reduced to scalar wave equations and the controlled hybridization approach [22] fails for 3D systems because of a singularity at Γ (cf. Sec. 3.4).

By expanding the Wannier functions in Eqn. (3.46) into their defining integrals and inserting Eqn. (3.47) we find:

$$\begin{aligned}
\int_{\mathrm{BZ}} \mathrm{d}k \sum_{ml} |\, H_{l\mathbf{k}}\,\rangle \mathcal{U}^{(\mathbf{k})}_{lm} \mathcal{D}^{(\hat{T},\mathrm{WF})}_{mn} &= \int_{\mathrm{BZ}} \mathrm{d}k \sum_{m} \hat{T} |\, H_{m\mathbf{k}}\,\rangle \mathcal{U}^{(\mathbf{k})}_{mn} && (3.48)\\
&= \int_{\mathrm{BZ}} \mathrm{d}k \sum_{ml} |\, H_{l\hat{T}\mathbf{k}}\,\rangle \mathcal{D}^{(\hat{T},\mathrm{BF},\mathbf{k})} \mathcal{U}^{(\mathbf{k})}_{mn} && (3.49)\\
&= \int_{\mathrm{BZ}} \mathrm{d}k \sum_{ml} |\, H_{l\mathbf{k}}\,\rangle \mathcal{D}^{(\hat{T},\mathrm{BF},\hat{T}^{-1}\mathbf{k})}_{lm} \mathcal{U}^{(\hat{T}^{-1}\mathbf{k})}_{mn}. && (3.50)
\end{aligned}$$

As the individual Bloch modes are orthogonal to each other, this must hold true for each **k**–point independently and we find:

$$\left[\mathcal{D}^{(\hat{T},\mathrm{BF},\hat{T}^{-1}\mathbf{k})} \mathcal{U}^{(\hat{T}^{-1}\mathbf{k})}_{mn}\right]^{-1} \mathcal{U}^{(\mathbf{k})}_{lm} \mathcal{D}^{(\hat{T},\mathrm{WF})}_{mn} = 1. \tag{3.51}$$

If this condition is satisfied for all operators $\hat{T}$ in the point group, the unitary mixing will yield Wannier functions that are hybridized according to Eqn. (3.46).

If we wish to follow the general idea of Sec. 3.3.6, we now have to define an auxiliary penalty term. Given the fact that all matrices on the left hand side of Eqn. (3.51) are unitary, the absolute value of the trace of this matrix product is always smaller or equal to N_b, where N_b is the number of bands. In fact, the trace can be N_b only if Eqn. (3.51) is fulfilled. Therefore, the simplest choice for a penalty term is

$$\Omega_{\mathrm{sym}} = \sum_{\hat{T}\in\mathcal{P}} \int_{\mathrm{BZ}} \mathrm{d}k \left(N_b - \mathrm{Tr}\left\{ [\mathcal{U}^{(\hat{T}^{-1}\mathbf{k})}]^{-1} [\mathcal{D}^{(\hat{T},\mathrm{BF},\hat{T}^{-1}\mathbf{k})}]^{-1} \mathcal{U}^{(\mathbf{k})} \mathcal{D}^{(\hat{T},\mathrm{WF})} \right\} \right). \tag{3.52}$$

The corresponding gradient reads (cf. Appendix A.2):

$$\begin{aligned}
\frac{\mathrm{d}\Omega_{\mathrm{sym}}}{\mathrm{d}\mathcal{W}^{(\mathbf{k})}} &= [\mathcal{U}^{(\mathbf{k})}]^{-1} [\mathcal{D}^{(\hat{T},\mathrm{BF},\mathbf{k})}]^{-1} \mathcal{U}^{(\hat{T}\mathbf{k})} \mathcal{D}^{(\hat{T},\mathrm{WF})}\\
&\quad - \mathcal{D}^{(\hat{T},\mathrm{WF})} [\mathcal{U}^{(\hat{T}^{-1}\mathbf{k})}]^{-1} [\mathcal{D}^{(\hat{T},\mathrm{BF},\hat{T}^{-1}\mathbf{k})}]^{-1} \mathcal{U}^{(\mathbf{k})}. && (3.53)
\end{aligned}$$

As the title of this section already suggests, I have not tried this, yet. However, I am very positive that it would work provided the problem of retrieving the Bloch modes' representations was solved. The main challenge could be caused by **k**–points with low symmetry such as the lines connecting high–symmetry points. Here, the Bloch modes' representations are not trivial (as they are for general **k**–points) but the **k**–points are not mapped onto themselves by every symmetry operator (as very high symmetry points such as Γ are). However, one might hope that if these points were excluded from the auxiliary penalty term (i. e. no symmetry would be directly imposed on them), the localization functional would be able to "interpolate" the symmetry condition if the neighboring **k**–points were forced to be symmetric (i. e. if they would contribute to the penalty term).

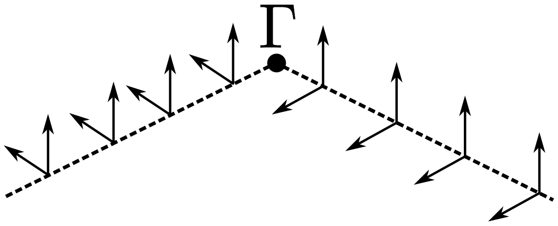

Figure 3.9.: *Illustration of the first kind of discontinuity caused by the divergence condition. The subspaces (pairs of arrows) for possible Fourier coefficients are invariant along each line of* $\mathbf{k}$*–vectors with the same direction (dotted lines). Consequently, there is a mismatch between continuations of subspaces along different directions to* Γ*. This necessarily leads to poorly localized Wannier functions for the first two bands.*

3.4. 3D Wannier Functions

Initially, my task as a PhD student was to generate Wannier functions for three–dimensional PhCs and to use them as a basis to simulate related defect structures such as cavities or wave guide structures. In this section, I present the results of this endeavor. Unfortunately, I was not able to construct functions that are sufficiently localized to serve as an efficient simulation basis.

3.4.1. Localization

In three dimensions, exponential localization of photonic Wannier function created from the physical Bloch modes of the crystal is impossible at least for some bands. This is due to fundamental topological discontinuities imposed by the absence of magnetic or electric charges [11], i. e. ,

$$\nabla \cdot \mathbf{H} = \nabla \cdot \mathbf{D} = 0. \tag{3.54}$$

In the remainder, I will stick to the magnetic formulation and comment on electric field Wannier functions in the end.

An exponentially localized generalized Wannier function is the lattice–Fourier transform of a band of generalized Bloch modes that is continuous throughout the whole Brillouin zone. This is fundamentally impossible for at least some bands of a three–dimensional photonic system. To show this we expand the Bloch functions $\mathbf{H}_{n\mathbf{r}}(\mathbf{k})$ into plane waves with vector-valued expansion coefficients $\widetilde{\mathbf{h}}_n(\mathbf{k}+\mathbf{G})$ (cf. Sec. 2.1.1). In this context, I regard the Bloch functions as functions of $\mathbf{k}$ with $\mathbf{r}$ being a parameter. This is due to the fact that we require continuity with respect to $\mathbf{k}$ for each real–space point

---

[11] The occurrence of magnetic charges would be in any case somewhat questionable.

independently. The expansion reads:

$$\mathbf{H}_{n\mathbf{k}}(\mathbf{r}) = \sum_{\mathbf{G}} \widetilde{\mathbf{h}}_n(\mathbf{k}+\mathbf{G}) \exp\big(\mathbf{i}\mathbf{r}\cdot(\mathbf{k}+\mathbf{G})\big)\,, \tag{3.55}$$

where the reciprocal lattice vectors are denoted by $\mathbf{G}$ and $\mathbf{k}$ is restricted to the first Brillouin zone. Since the exponential is holomorphic in $\mathbf{K} = \mathbf{k}+\mathbf{G}$, any discontinuity must stem from $\mathbf{h}_n(\mathbf{K})$.

The divergence condition is easily expressed in the plane wave basis:

$$\begin{aligned} \mathbf{K}\cdot\widetilde{\mathbf{h}}(\mathbf{K}) &= 0, & (3.56)\\ |\widetilde{\mathbf{h}}(\mathbf{K})| &= 1. & (3.57) \end{aligned}$$

This condition is well–known from the classic plane wave band structure computation (Sec. 2.1) where it is used to reduce the problem size by one third. However, it causes two types of discontinuities. The first one is relevant for the lowest bands only, whereas the second, more subtle type, in principle may affect every band.

First, it is worth noting that Eqn. (3.56) spans a two–dimensional solution space for $\widetilde{\mathbf{h}}(\mathbf{K}\neq\mathbf{0})$. At the Γ–point ($\mathbf{K}=\mathbf{0}$) every vector is possible. The physical interpretation is that there are only two polarizations for propagating light waves, whereas static fields may be oriented arbitrarily. Aiming at a continuous set of polarization vectors throughout the BZ, we must restrict ourselves to a two–dimensional subspace at Γ. For a fixed $\mathbf{K}_0$, the solution space is invariant under a scaling transformation $\mathbf{K} = \lambda\mathbf{K}_0$ with $\lambda \neq 0$. By taking the limit $\lambda \to 0$, we can construct a two–dimensional subspace at Γ such that the solution spaces along the direction from the Γ–point to $\mathbf{K}_0$ are continuous. However, this is the *only* direction along which the set of function spaces is continuous. This first discontinuity is a discontinuity of the polarization spaces as a function of $\mathbf{K}$ near the Γ–point (cf. Fig. 3.9).

The second issue is closely related to the problem of combing a hairy sphere and arises due to Eqn. (3.57). We consider a sphere around Γ with radius $|\mathbf{K}|$. The solution space at each point is given by the tangential plane to the sphere. Starting with a given basis at a point $\mathbf{K}$, we approach the point $-\mathbf{K}$ on a great circle of the sphere and return to $\mathbf{K}$ on another great circle (cf. Fig. 3.10). On each path, we translate the basis adiabatically. The limit for approaching $\mathbf{K}$ differs from the initial orientation[12]. This is just a simplified picture to demonstrate the general problem. As can be shown rigorously [46], every continuous vector–valued function that is defined on a sphere and fulfills Eqn. (3.56) has to vanish at least at one point. Conversely, a non–vanishing function thus has at least one point of discontinuity, although the tangential spaces of the sphere are continuous with respect to $\mathbf{K}$.

The above discussion also applies to distorted spheres (such as cubes or icosahedra) as long as they enclose Γ. The complete reciprocal space is tessellated by reciprocal unit cells, each of which is centered around a reciprocal lattice vector $\mathbf{G}$. Hence, each point

---

[12] This is also a popular example for the fact that rotations around different axes do not commute.

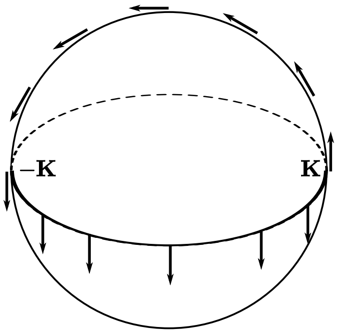

Figure 3.10.: *Sketch to motivate the general problem leading to the second discontinuity that arises due to the divergence condition. If a vector (e.g. a basis vector of the tangential subspaces) is adiabatically transferred from* **K** *to* −**K** *around the upper path and returned via the horizontal path, it gains a minus sign. Of course, a continuous set of vectors can be found along the sketched path.*

$\mathbf{K}$ is related to one (the closest) $\mathbf{G}$–vector and a sphere situated in the first Brillouin zone can contain the Γ–point if and only if $\mathbf{G} = \mathbf{0}$[13]. Therefore, the discontinuities occur for $\widetilde{\mathbf{h}}(\mathbf{k} + \mathbf{0})$ only.

Due to the second type discontinuity, every band that contains a non–vanishing contribution of $\widetilde{\mathbf{h}}(\mathbf{k} + \mathbf{0})$ on some subspace Ω_{sub} of the first BZ is inherently discontinuous with respect to $\mathbf{k}$ as soon as Ω_{sub} supports a surface enclosing Γ. This is necessarily true for band one and two because of the effective-medium limit in the vicinity of the Γ–point. However, due to the periodic scattering inside a photonic crystal with finite index contrast, higher bands contain $\mathbf{G} = \mathbf{0}$ amplitudes, as well. For crystals with a step-discontinuous dielectric function, the contamination is roughly reciprocal to the band's frequency. Whether higher bands are fundamentally discontinuous depends on the sym-

---

[13] For each point $\mathbf{K}_0$, we can construct two basis functions $\widetilde{\mathbf{h}}_{1,2}(\mathbf{K})$ that obey Eqns. (3.56, 3.57) and are continuous in the open half–space. $\mathbf{K} \cdot \mathbf{K}_0 > 0$:

$$\begin{aligned}
\widetilde{\mathbf{h}}_1(\mathbf{K}) =& \frac{\mathbf{v}(\mathbf{K})}{|\mathbf{v}(\mathbf{K})|}, \\
\widetilde{\mathbf{h}}_2(\mathbf{K}) =& \frac{\mathbf{K} \times \mathbf{v}(\mathbf{K})}{|\mathbf{K} \times \mathbf{v}(\mathbf{K})|}, \qquad (3.58) \\
\text{where} \quad \mathbf{v}(\mathbf{K}) =& \mathbf{K} \times \mathbf{K} \times \widetilde{\mathbf{h}}(\mathbf{K}_0).
\end{aligned}$$

Please note that $\mathbf{v}(\mathbf{K})$ is continuous in the whole space but vanishes on the axis $\mathbf{K}_{\text{van}} = \lambda \widetilde{\mathbf{h}}(\mathbf{K}_0)$ and that the limit

$$\lim_{\mathbf{K} \to \mathbf{K}_{\text{van}}} \mathbf{v}(\mathbf{K})/|\mathbf{v}(\mathbf{K})|$$

does not exist. In particular, the functions defined by Eqn. (3.58) are continuous on each sphere that lies completely in one half–space. For a given sphere not enclosing Γ, this holds for instance if we choose $\mathbf{K}_0$ as the sphere's center. Hence, only spheres enclosing Γ are subject to the second type discontinuity.

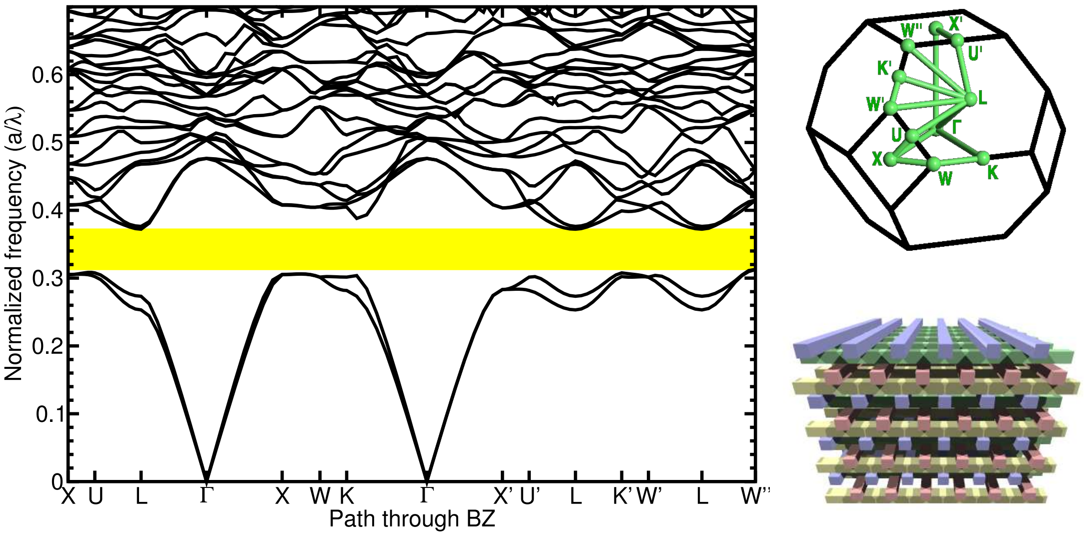

Figure 3.11.: *The band structure of the an fcc–woodpile (left panel, square rods,* $\epsilon = 12$*), first Brillouin zone with the band structure path (top right panel) and sketch of the physical structure (bottom right, colors were added to improve contrast and have no physical meaning). The photonic band gap is highlighted in the band structure plot in yellow.*

metry of the dielectric distribution—more precisely whether the $(\mathbf{G} = \mathbf{0})$–contributions support a closed surface or form isolated areas in the first Brillouin zone.

We have seen that the eigenmodes of two bands in a 3D photonic crystal must be discontinuous. Furthermore, I motivated why in fact no band can safely be assumed to provide a **k**–continuous basis. However, the Fourier transform of a discontinuous function has an envelope decaying with at best first order localization $(\mathcal{O}(R^{-1}))$. Thus, we may not expect a well–localized magnetic Wannier basis for 3D photonic crystals.

The discussion above can be identically formulated based on electric displacement fields **D**. They are connected to the electric fields via the constitutive relation

$$\mathbf{D}(\mathbf{r}) = \epsilon(\mathbf{r})\mathbf{E}(\mathbf{r}). \tag{3.59}$$

Since the position **r** is a parameter and ϵ does not depend on **k**, the above results directly apply to the electric field formulation. Furthermore, all results also apply to dispersive and anisotropic systems.

3.4.2. Woodpile Wannier Functions

Finally, I did succeed in a creating a set of 71 Wannier functions for an fcc–woodpile structure of silicon rods (similar to [47], however with $\epsilon = 12$ and rods with square cross section). I constructed these modes from a $7 \times 7 \times 7$–discretization of the first Brillouin zone at a spatial resolution of $32 \times 32 \times 32$ discretization points per unit cell.

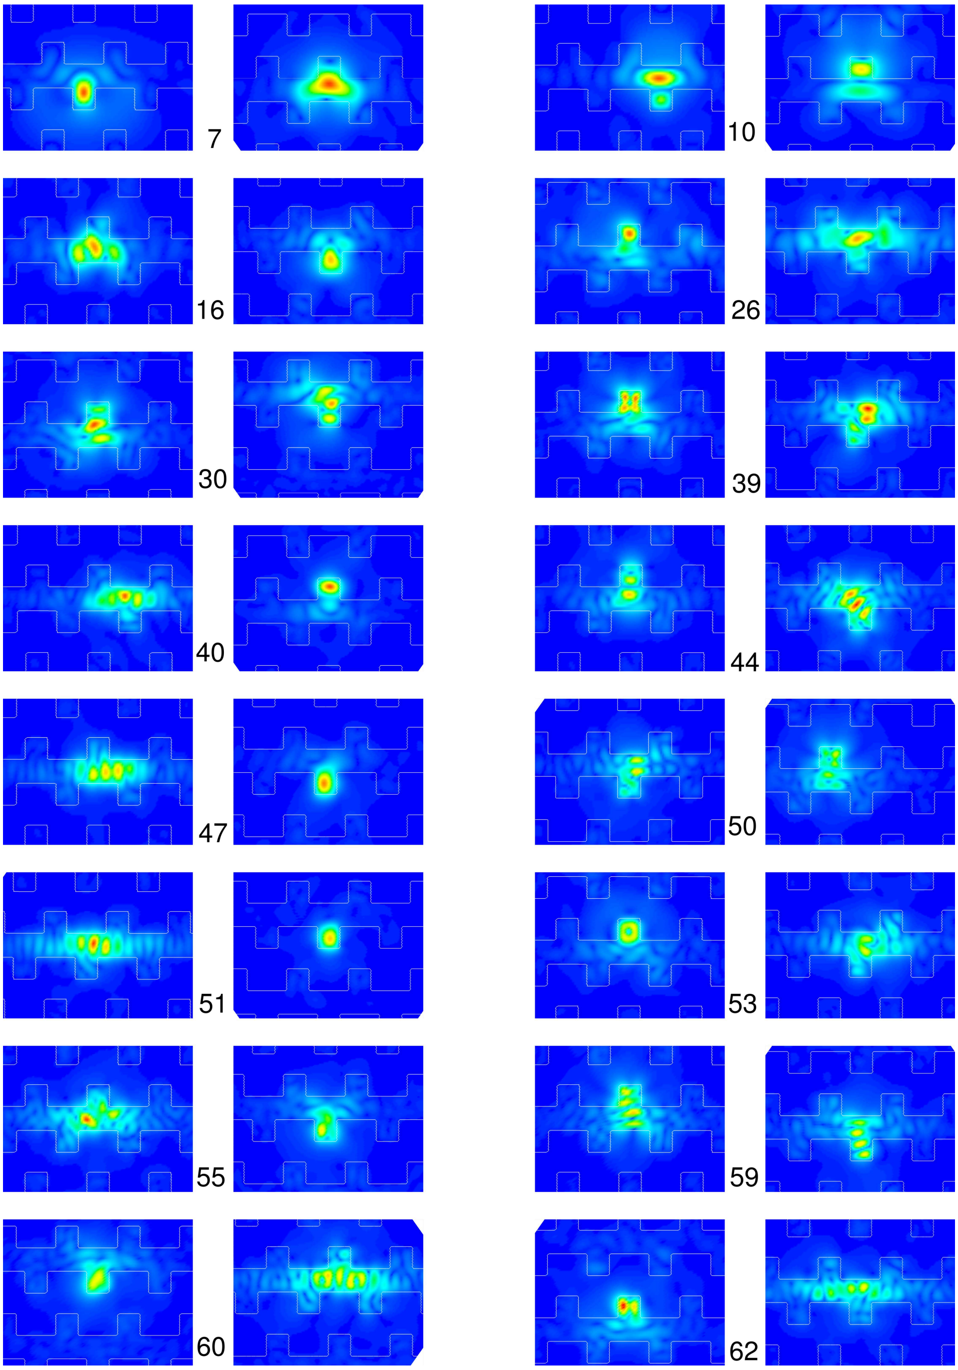

Figure 3.12.: *Plots of some 3D Wannier functions for the woodpile structure (fcc–lattice, square rods, $\epsilon = 12$). The numbers indicate the band numbers, the surrounding plots show the absolute value of the magnetic field amplitude of the WF cut along two orthogonal planes through (or close to) the intensity maximum. The outline of the dielectric structure is shown as white lines.*

For minimizing the spread, I employed the conjugate gradients method in combination with the simulated annealing super–minimizer for higher bands. Furthermore, I reduced the imaginary parts using the auxiliary penalty approach presented in Sec. 3.3.6.

The respective band structure together with a sketch of the system is depicted in Fig. 3.11. The first two bands have already been presented in Fig. 3.8. In Fig. 3.12, I show some of the higher bands. As compared to the first two bands (which strongly suffer from the discontinuity discussed above), some of them tend to be more symmetric and most of them are better localized. However, all functions except for the first few are rather distorted, some of them such as band 50 or band 55 are completely erratic. However, even a single poorly localized function would be already sufficient to drastically reduce the efficiency of the Wannier method. As a consequence, the photonic Wannier basis unfortunately turned out not to form an efficient basis in 3D.

The **k**–space resolution that I used to create this function set is rather coarse. Unfortunately, this already required several dozen gigabyte of hard disk space just for the Bloch functions and resulted in a minimization process with over 300000 degrees of freedom for the higher band groups. For increasing frequencies, the band groups become larger because the spectral density of bands grows roughly quadratically with increasing frequency whereas the size of the frequency windows that are required for artificially isolating the band group (cf. Sec. 3.2.2) roughly remains constant. I created Wannier functions from finer discretizations up to $13 \times 13 \times 13$ **k**–points as well, but was unable to get past some thirty functions before the minimization problem ceased to be manageable. Furthermore, the localization properties of the resulting WFs were not superior to those presented in Fig. 3.12.

3.5. Wannier–Based Simulation

Although photonic Wannier functions for three–dimensional systems are not well–localized, the corresponding functions for two–dimensional crystals have proven to form an efficient expansion basis for the simulation of very large PhC–based systems, e. g. , via a scattering matrix formalism [31]. In particular, the E–polarization features very fast convergence of defect mode frequencies and mode patterns. In such systems, acceptable results can be obtained with as few as eight degrees of freedom per unit cell. Thus, the Wannier basis is a good choice for investigating numerically challenging fundamental properties of PhC–structures. One such example is the systematic investigation of the consequences of disorder in photonic crystals. During his time as a diploma student in our group, Martin Köhl performed a full analysis of several approaches to this problem [14, 15].

3.5.1. Discretization

Within the Wannier method, all fields are expanded into the basis of WFs:

$$\mathbf{H}(\mathbf{r}) = \sum_{\mathbf{R}} \sum_{n} c_{n\mathbf{R}}^{(H)} \mathbf{W}_{n\mathbf{R}}(\mathbf{r}), \tag{3.60}$$

where $c_{n\mathbf{R}}^{(H)}$ may be real–valued (for time–domain simulations) or complex–valued (for the frequency–domain) expansion coefficients. The first sum covers all lattice sites inside a computational domain. The second sum runs over all considered bands, where each lattice site may be assigned a different number of basis functions. This allows to discretize individual lattice with different accuracy.

Next, we expand one of the wave equations in Sec. 1.1 into our basis. Here, I illustrate this starting with Eqn. (1.23). First, we expand the fields according to Eqn. (3.60). Next, we project the whole equation onto a set of test functions, which we choose to be the expansion basis itself (Galerkin choice). From this we get the discretized wave equation in form of a matrix problem:

$$-\frac{1}{c^2} \sum_{n'\mathbf{R}'} \underbrace{\langle\, \mathbf{W}_{n\mathbf{R}} \,|\, \mathbf{W}_{n'\mathbf{R}'} \,\rangle}_{\mathcal{A}^{(H)}_{nn',\mathbf{R}\mathbf{R}'}} \partial_t^2 c_{n'\mathbf{R}'}^{(H)} = \sum_{n'\mathbf{R}'} \underbrace{\langle\, \mathbf{W}_{n\mathbf{R}} \,|\nabla \times (\eta \nabla \times)|\, \mathbf{W}_{n'\mathbf{R}'} \,\rangle}_{\mathcal{M}^{(H)}_{nn',\mathbf{R}\mathbf{R}'}} c_{n'\mathbf{R}'}^{(H)}. \tag{3.61}$$

Here, $\eta(\mathbf{r}) = \epsilon^{-1}(\mathbf{r})$ denotes the inverse dielectric function and the superscript (H) indicates that the matrices were derived within the magnetic–field formulation of the Wannier approach. The speed of light was set to one (cf. Eqn. (1.27) through Eqn. (1.30)).

Typically, our Wannier–based simulation tools are applied to systems where a host photonic crystal is only locally modified, i. e. contains defects. Therefore, it is convenient to decompose the material distribution $\eta(\mathbf{r})$ into the contribution $\eta^{(\mathrm{per})}(\mathbf{r})$ from the ideal PhC backbone and the local deviations $\delta\eta(\mathbf{r})$:

$$\eta(\mathbf{r}) = \eta^{(\mathrm{per})}(\mathbf{r}) + \delta\eta(\mathbf{r}), \tag{3.62}$$

$$\eta^{(\mathrm{per})}(\mathbf{r}) = \eta^{(\mathrm{per})}(\mathbf{r} + \mathbf{R}). \tag{3.63}$$

Accordingly, the system matrix consists of two contributions:

$$\mathcal{M}^{(H)} = \mathcal{C}^{(H)} + \mathcal{D}^{(H)}, m \tag{3.64}$$

$$\mathcal{C}^{(H)}_{nn',\mathbf{RR}'} := \langle \nabla\times\mathbf{W}_{n\mathbf{R}} \,|\eta^{(\mathrm{per})}|\, \nabla\times\mathbf{W}_{n'\mathbf{R}'} \rangle, \tag{3.65}$$

$$\mathcal{D}^{(H)}_{nn',\mathbf{RR}'} := \langle \nabla\times\mathbf{W}_{n\mathbf{R}} \,|\delta\eta|\, \nabla\times\mathbf{W}_{n'\mathbf{R}'} \rangle, \tag{3.66}$$

where $\mathcal{C}$ is block–diagonal in the Wannier basis, sparse and contains the full information about the photonic band structure. The defect matrix $\mathcal{D}$ on the other can be decomposed into contributions that stem from the individual point–like defects [30]:

$$\mathcal{D}^{(H)} = \sum_{i=1}^{N^{(\mathrm{def})}} b_i \mathcal{B}_{nn',(\mathbf{R}-\mathbf{R}_i^{(\mathrm{def})})(\mathbf{R}'-\mathbf{R}_i^{(\mathrm{def})})}. \tag{3.67}$$

Here, $N^{(\mathrm{def})}$ is the number of defects, $\mathbf{R}_i^{(\mathrm{def})}$ is the position of the i–th defect and b_i is a scattering strength, which corresponds to the index contrast of the defect. The atomic defect matrix $\mathcal{B}$ is defined analogously to Eqn. (3.66).

We end up with the semi–discrete form of the wave equation:

$$\mathcal{A}^{(H)}\partial_t^2 \vec{c}^{\,(H)}(t) = -(\mathcal{C}^{(H)} + \mathcal{D}^{(H)})\vec{c}^{\,(H)}(t). \tag{3.68}$$

The name refers to the fact that all spatial derivatives have been eliminated by the field discretization Eqn. (3.60) but the temporal derivatives are still present. Consequently, we have reduced the problem to a vector–valued ordinary differential equation with respect to t.

An analogous semi–discrete equation can be derived for the electric field formulation starting with Eqn. (1.24):

$$(\mathcal{C}^{(E)} + \mathcal{D}^{(E)})\partial_t^2 \vec{c}^{\,(E)}(t) = -\mathcal{A}^{(E)}\vec{c}^{\,(E)}(t). \tag{3.69}$$

Here, the coefficient matrices are slightly differently defined:

$$\mathcal{A}^{(E)} := \langle \nabla\times\mathbf{W}_{n\mathbf{R}} \,|\, \nabla\times\mathbf{W}_{n'\mathbf{R}'} \rangle, \tag{3.70}$$

$$\mathcal{C}^{(E)} := \langle \mathbf{W}_{n\mathbf{R}} \,|\epsilon^{(\mathrm{per})}|\, \mathbf{W}_{n'\mathbf{R}'} \rangle, \tag{3.71}$$

$$\mathcal{D}^{(E)} := \langle \mathbf{W}_{n\mathbf{R}} \,|\delta\epsilon|\, \mathbf{W}_{n'\mathbf{R}'} \rangle, \tag{3.72}$$

$$\epsilon(\mathbf{r}) = \epsilon(\mathbf{r})^{(\mathrm{per})} + \delta\epsilon(\mathbf{r}). \tag{3.73}$$

3.5.2. Frequency–Domain Methods

The semi–discrete equations (3.68) and (3.69) can be solved for time–harmonic solutions by performing a Fourier transform. Effectively, this corresponds to replacing the temporal derivative with the frequency:

$$\partial_t \to -\mathrm{i}\omega. \tag{3.74}$$

We obtain the fully discretized matrix equation

$$\sum_{n',\mathbf{R}'} \left[\mathcal{C}_{nn',\mathbf{RR}'} + \mathcal{D}_{nn',\mathbf{RR}'} - \lambda\mathcal{A}_{nn',\mathbf{RR}'}\right] c_{n'\mathbf{R}'} = \vec{s}, \tag{3.75}$$

where $\vec{s}$ represents a source term and the factor λ contains the frequency ω, depending on whether the magnetic or the electric formulation was chosen:

$$\lambda^{(E)} = \frac{1}{\omega^2}, \tag{3.76}$$

$$\lambda^{(H)} = \omega^2. \tag{3.77}$$

There are two main approaches to solve Eqn. (3.75):

1. The source term can be set to zero and the equation can be regarded as an eigenvalue problem with respect to λ. This yields the eigenmodes and eigenfrequencies of the system under consideration. Typical problems are the calculation of cavity resonances.

2. The frequency parameter λ can be regarded as a parameter. In this case, the system typically is a wave guiding structure such as an integrated interferometer [31]. The source term is used to launch a wave guide mode and the calculation results in transmission and reflection spectra of the photonic device.

These two methods as well as a wave guide solver were developed by Daniel Hermann and presented in detail in his PhD thesis [30].

3.5.3. Time–Domain Method

A time–domain simulation software based on photonic Wannier functions was developed to complement our frequency–domain numerics. This work was performed by Christian Blum as the topic of his diploma thesis [48, 49] under my supervision and enabled by my numerical framework. In this section, I review the main part of his thesis.

So far, this method has only been implemented for E–polarization in one– and two–dimensional systems. We start with the respective semi–discrete wave equation Eqn. (3.69) which is an implicit ordinary differential equation of second order in t:

$$\mathcal{A}^{(E)}\vec{c}^{\,(E)} + (\mathcal{C}^{(E)} + \mathcal{D}^{(E)})\partial_t^2\vec{c}^{\,(E)}, \tag{3.78}$$

where we only consider real–valued expansion coefficients.

In principle, Eqn. (3.78) could be recast into an explicit differential equation (i. e. where the temporal derivative and the matrix–vector product are separated) by multiplying the whole equation with $(\mathcal{C} + \mathcal{D})^{-1}$. However, the time required for calculating this matrix inverse scales at least quadratically[14] with the number of degrees of freedom. This would drastically limit the manageable system size.

---

[14] It scales quadratically if the sparsity of the matrix is exploited; otherwise it scales cubically.

This decision to solve the wave equation as an implicit problem reduces the time integrators available to evolve the photonic state. For instance, it excludes most Runge–Kutta–type methods. Before we can discuss the time integration, however, it is necessary to reduce the order of the differential equation by adding the first temporal derivative of the Wannier expansion coefficients as an auxiliary set to the degrees of freedom:

$$\underbrace{\begin{pmatrix} 1 & 0 \\ 0 & (\mathcal{C}+\mathcal{D}) \end{pmatrix}}_{\mathcal{M}} \partial_t \underbrace{\begin{pmatrix} \vec{c} \\ \vec{d} \end{pmatrix}}_{\vec{\psi}} = \underbrace{\begin{pmatrix} 0 & 1 \\ -\mathcal{A} & 0 \end{pmatrix}}_{\mathcal{S}} \begin{pmatrix} \vec{c} \\ \vec{d} \end{pmatrix}. \tag{3.79}$$

The matrices in this first order implicit equation are commonly referred to as the mass matrix $\mathcal{M}$ and the stiffness matrix $\mathcal{S}$[15]. Next, the temporal derivative is approximated by backward differences where the time evolution is assumed to be done for equidistant time steps of width Δt:

$$\Delta t \partial_t \vec{\psi}(t_0) = \sum_{i=0}^{\infty} \alpha_i \vec{\psi}(t_0 - i\Delta t). \tag{3.80}$$

In practice, the summation must be truncated to a finite number N of summands. The values for the difference coefficients α_i for various difference orders (i. e. numbers of summands) can be found in Ref. [48]. However, we have only used second order differences in practice.

By inserting Eqn. (3.80) into Eqn. (3.79) and collecting all terms for time t_0 on the left hand side, we obtain a system of linear equations:

$$(\Delta t \mathcal{S} - \alpha_0 \mathcal{M})\psi(t_0) = \sum_{i=1}^{N} \alpha_i \mathcal{M} \psi(t_0 - i\Delta i). \tag{3.81}$$

This equation allows to evolve a state in time but requires to solve a system of linear equations in each time step. However, the matrix $(\Delta t\mathcal{S} - \alpha_0\mathcal{M})$ is sparse and does not change during the simulation. Thus, we can perform a sparse LU–decomposition and perform a back–substitution in each iteration. This is very fast for small systems but scales quadratically with respect to time as the system size is increased. For larger systems, Krylov–subspace solvers similar to the conjugate gradient method mentioned in Sec. 3.3.3 perform favorably. As a starting point for the minimization, we use the Wannier coefficients at the previous time step hoping that the photonic state does not change too much within the time interval Δt. Indeed, the solver finds the solution within very few iterations and we observe essentially linear scaling between execution time and system size [49].

The real strength of time–domain methods is their fundamental ability to simulate nonlinear systems, which are very hard to handle in frequency–domain. This is possible

---

[15] These terms stem from continuum mechanics which was one of the first fields to employ finite element formulations.

in the Wannier framework as well. It was proposed by Christian Blum and demonstrated by Alexandra M. Graham during her time as a master student under my supervision.

In a nutshell, the only change required to handle nonlinear systems is a nonlinear coupling between the expansion coefficients which has to be added to the wave equation. In the case of a defect with $\chi^{(2)}$–nonlinearity, this means that the corresponding “overlap coefficient matrix” $\mathcal{Q}^{(2)}$ has to be an object that yields an expansion coefficient vector if multiplied with two state vectors instead of only one as in the case of a linear defect. Therefore, it has to have three pairs of band and lattice site indices and is the overlap integral of three Wannier functions and the defect kernel:

$$\mathcal{Q}^{(2)}_{nn'n'',\mathbf{R}\mathbf{R}'\mathbf{R}''} = \int \mathrm{d}r \; \mathbf{W}^{*}_{n\mathbf{R}}(\mathbf{r}) \; \chi^{(2)}(\mathbf{r}) \; \mathbf{W}_{n'\mathbf{R}'}(\mathbf{r}) \; \mathbf{W}_{n''\mathbf{R}''}(\mathbf{r}), \tag{3.82}$$

and the nonlinear semi–discrete wave equation simply becomes

$$\begin{aligned}\sum_{n'\mathbf{R}'} \Big[\mathcal{C}_{nn',\mathbf{R}\mathbf{R}'} + \mathcal{D}_{nn',\mathbf{R}\mathbf{R}'} + \Big(\sum_{n''\mathbf{R}''} \mathcal{Q}_{nn'n'',\mathbf{R}\mathbf{R}'\mathbf{R}''} \; \partial_t d_{n''\mathbf{R}''}\Big)\Big] \partial_t d_{n'\mathbf{R}'} \\ = -\sum_{n'\mathbf{R}'} \mathcal{A}_{nn',\mathbf{R}\mathbf{R}'} c_{n'\mathbf{R}'}.\end{aligned} \tag{3.83}$$

One of the major difficulties in this approach in higher dimensions is the vast number of overlap elements in $\mathcal{Q}^{(2)}$. Alexandra Graham implemented and tested this nonlinear extension for one–dimensional systems in the electric field formulation [50].

3.5.4. Absorbing Boundary Conditions

Absorbing boundary conditions are a crucial component for the simulation of open systems, i. e. of systems that can be left by a wave without reflections. The outgoing radiation must leave a core area of the simulation domain and never return. In systems surrounded by a homogeneous material (such as vacuum), this is usually accomplished with perfectly matched layers (PML) [9]. Those are areas filled with an unphysical material that is highly dissipative but has the same impedance as the material to be extended. Consequently, the wave traverses the interface to the PML–region without reflection and is then attenuated by several orders of magnitude before it is eventually reflected at the far side of the PML.

Unfortunately, this approach fails inside a photonic crystal [51]. The PML interface becomes highly reflective. This can be attributed to the fact that “impedance” becomes spatially dependent in such a structured medium (or that it is no longer a well–defined quantity after all). Actually, phase and group velocity as well as the field patterns of the eigenmodes have to be matched between the regular simulation domain and the absorbing boundary region. In fact, perfectly matched layers are no longer superior to a simple adiabatic absorber [48, 49] inside a PhC.

However, the typical application of the Wannier method (photonic circuits with wave guides operated in the band gap of the host crystal) restricts the problem to terminating one–dimensional wave guides.

In this section, I discuss several approaches to accomplish this both in time–domain and in frequency–domain. First, I discuss the differences between these two problem classes. Then, I review the technique of adiabatic absorbers, which is our standard absorber for the Wannier function time–domain simulator and was implemented and tested by Christian Blum [48]. Furthermore, I review the basis transformation boundary condition, which was developed by Daniel Hermann [30] and is our standard boundary condition for frequency–domain scattering matrix problems. However, I provide an alternative explanation for its function. This understanding allows to reduce the computational effort as presented in Sec. 3.5.4. In the remainder of this section, I discuss two other approaches which I considered predominantly for developing absorbing boundaries in the time–domain. Unfortunately, none of the two last methods exhibits remotely acceptable performance.

Frequency–Domain vs. Time–Domain

In principle, absorbing boundary conditions face similar requirements both for time–domain and frequency–domain simulations: They are supposed to allow radiation to reflectionlessly leave the core simulation domain and prevent it from returning once it left. However, there are different possible approaches and for each approach may be incorporated differently in the master equations. In particular, the time–domain and the frequency–domain formulation impose different typical restrictions.

Lossy materials are a prerequisite for some absorbing boundary approaches (such as PMLs and adiabatic absorbers) that aim to suppress reflections by dissipatively attenuating the outgoing wave before it hits the (reflecting) boundary of the computational domain. Losses can be added by introducing a damping term $\mathcal{D}'\partial_t c$ into the wave equation, where $\mathcal{D}'$ describes an electric conductance defect and is defined in analogy to Eqn. (3.66). The lossy semi–discrete wave equation then reads:

$$(\mathcal{C}^{(E)} + \mathcal{D}^{(E)})\partial_t^2 \vec{c}^{\,(E)}(t) + \mathcal{D}'^{(E)}\partial_t \vec{c}^{\,(E)}(t) + \mathcal{A}^{(E)}\vec{c}^{\,(E)}(t) = 0. \tag{3.84}$$

In time–domain calculations, this directly results in additional terms in the linear systems of equations [48]. In the frequency–domain framework, the same effect can be accomplished to adding an imaginary part to the dielectric defect distribution in Eqn. (3.66) or Eqn. (3.72), respectively. As the expansion coefficients may be complex numbers in a frequency–domain calculation, no changes have to be made to the loss–less implementation.

An example for a restriction is the handling of frequency–dependent quantities such as a dispersive dielectric constant. This can be trivially handled in frequency–domain problems where the frequency enters as a fixed parameter: The dielectric constant is simply readjusted for each frequency. In time–domain, on the other hand, it generally results in a computationally very expensive convolution[16].

---

[16] This can be avoided by means of auxiliary fields for some special types of dispersive behavior such as Drude–type materials [9].

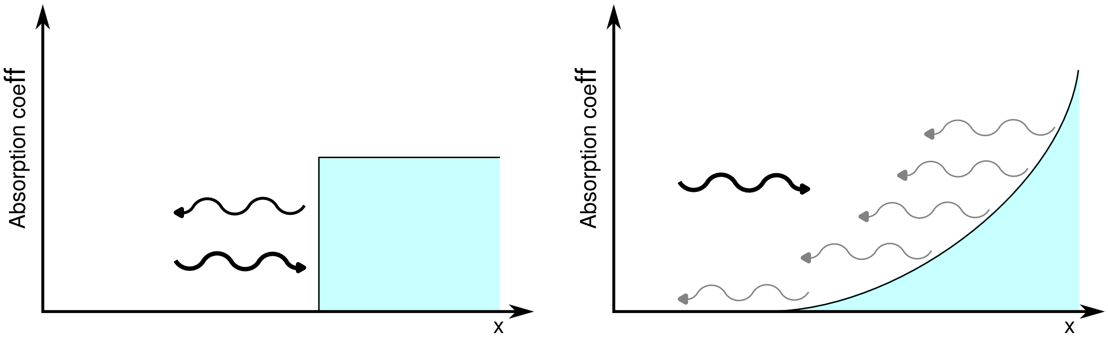

Figure 3.13.: *Reflections at the transition between a lossless domain and a lossy domain. If the transition is abrupt, this interface may cause strong reflections. In a material whose absorptivity is spatially slowly increased, however, there is a distributed reflection that decreases as the absorption profile is flattened.*

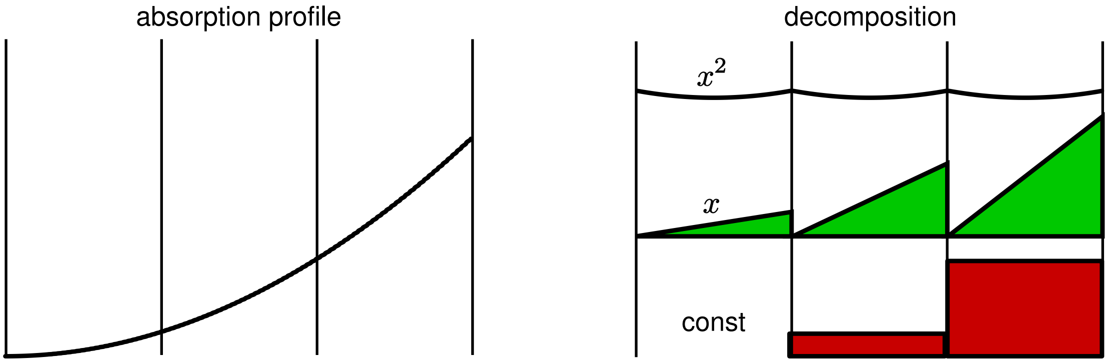

Figure 3.14.: *Sketch of how the absorption profile (here: second order parabola) of an adiabatic absorber that is extended over several (here: three) unit cells can be decomposed into monomials on each unit cell.*

Adiabatic Absorbers

An adiabatic absorber aims to dissipate the entering wave. Reflections at the boundary between the regular simulation domain and the absorber are reduced by slowly ramping up the absorption coefficient [51]. However, a distributed reflection of the wave due to the continuously varying absorptivity is unavoidable (see Fig. 3.13). This approach works fairly well both in time–domain as well as in frequency–domain.

It can be implemented efficiently by decomposing the whole (several unit cells long) absorber into pieces each of which is one unit cell in size. On each unit cell, the absorption profile is expanded in a polynomial basis (cf. Fig. 3.14) where each basis function corresponds to one dissipative defect matrix. This also allows to approximate transcendent profiles such as exponentials by spline interpolation. In principle, every class of

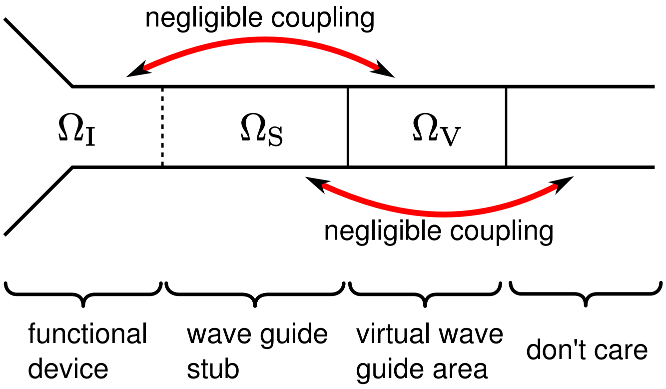

Figure 3.15.: *Decomposition of the simulated system into three parts: The functional device area contains the circuit to be simulated and which is connected to (in theory infinitely long) wave guides. Each device terminal is connected to a wave guide stub which is not part of the device itself but required to launch incoming wave guide modes and to terminate outgoing radiation. Finally, the wave guide stub is virtually extended. This area is only introduced to derive the absorbing boundary condition and does not appear in the numerics, hence the name "virtual wave guide area".*

basis polynomials could be used. For simplicity we chose monomials, which have the additional advantage that the polynomial order of the absorber can be increased by simply adding higher order dissipative defect matrices. In each absorber unit cell, the dissipative defect strength factors (corresponding to the factors b_i in Eqn. (3.67)) are then identical to the polynomial coefficients of the spline.

Basis Transformation Absorbers

Daniel Hermann developed and implemented a type of absorbing boundaries within his scattering matrix framework [30, 31]. This approach is based on expanding the fields in a wave guide stub into the basis of wave guide eigenmodes. In the formulation presented in this chapter, it is only suitable for frequency–domain simulations.

We start with the general form Eqn. (3.75) of the frequency–domain wave equation:

$$\underbrace{\left[\mathcal{C} + \mathcal{D} - \lambda\mathcal{A}\right]}_{\mathcal{L}} \vec{\psi} = \vec{s}, \tag{3.85}$$

where $\mathcal{L}$ is a linear operator expanded into the Wannier basis and λ is regarded as a fixed parameter.

Next, we decompose the lattice into three sub–domains. The inner domain Ω_I contains the circuit to be simulated. As it is assumed to be operated at frequencies within the band gap of the host crystal, radiation can only enter or leave via wave guides connected at dedicated ports. As these wave guides theoretically are infinitely long, the ports must be connected to a combination of absorbing boundary and source.

Both requires a wave guide stub Ω_{S} to be connected to each port of the device. Finally, another piece of the wave guide is required next to the boundary/source–stub in order to derive the correct boundary condition. I would like to call this the virtual wave guide area Ω_{V}. This decomposition is illustrated in Fig. 3.15 for a single wave guide port.

The additional areas Ω_{S} and Ω_{V} must be chosen so wide that the coupling between Ω_{I} and Ω_{V} as well as the coupling between Ω_{S} and the wave guide beyond Ω_{V} are negligible. Consequently, we may decompose the system matrix into the "on–domain" blocks $\mathcal{L}_{\mathrm{II}}$, $\mathcal{L}_{\mathrm{SS}}$, $\mathcal{L}_{\mathrm{VV}}$ and the coupling blocks $\mathcal{L}_{\mathrm{IS}}$, $\mathcal{L}_{\mathrm{SI}}$, $\mathcal{L}_{\mathrm{SV}}$ and $\mathcal{L}_{\mathrm{VS}}$:

$$\begin{pmatrix} \mathcal{L}_{\mathrm{II}} & \mathcal{L}_{\mathrm{IS}} & 0 \\ \mathcal{L}_{\mathrm{SI}} & \mathcal{L}_{\mathrm{SS}} & \mathcal{L}_{\mathrm{SV}} \\ 0 & \mathcal{L}_{\mathrm{VS}} & \mathcal{L}_{\mathrm{VV}} \end{pmatrix} \cdot \begin{pmatrix} \vec{\psi}_{\mathrm{I}} \\ \vec{\psi}_{\mathrm{S}} \\ \vec{\psi}_{\mathrm{V}} \end{pmatrix} = \begin{pmatrix} \vec{s}_{\mathrm{I}} \\ \vec{s}_{\mathrm{S}} \\ 0 \end{pmatrix}. \tag{3.86}$$

Accordingly, we decomposed the state vector and the source terms into contributions from the three domains as well.

If every outward directed mode in Ω_{S} is correctly continued in Ω_{V}, no reflections at the interface between both domains occur. On the other hand, both domains form adjacent sections of the same wave guide. Thus, we can continue the fields in Ω_{S} via the wave guide's propagator $\mathcal{P}_{\mathrm{VS}}$:

$$\vec{\psi}_{\mathrm{V}} = \mathcal{P}_{\mathrm{VS}} \vec{\psi}_{\mathrm{S}}, \tag{3.87}$$

where $\mathcal{P}_{\mathrm{VS}}$ extends all eigenmodes in the stub area to the virtual area according to their propagation constant k:

$$\mathcal{P}_{\mathrm{VS}} = \mathcal{E}_{\mathrm{WG}} \begin{pmatrix} \exp\big(\mathrm{i}k_1 \Delta R\big) & & 0 \\ & \ddots & \\ 0 & & \exp\big(\mathrm{i}k_N \Delta R\big) \end{pmatrix} \mathcal{E}_{\mathrm{WG}}^{-1}. \tag{3.88}$$

Here, $\mathcal{E}_{\mathrm{WG}}$ is the matrix that diagonalizes the wave guide problem for the given frequency and is unitary for a lossless wave guide. It consists of the Wannier coefficients of all (outward propagating and outward decaying evanescent) eigenmodes and I would like to refer to it as the eigenmode matrix; ΔR is the width of the wave guide stub Ω_{S} measured in lattice sites.

By inserting Eqn. (3.87) into Eqn. (3.86), we find

$$\begin{pmatrix} \mathcal{L}_{\mathrm{II}} & \mathcal{L}_{\mathrm{IS}} & 0 \\ \mathcal{L}_{\mathrm{SI}} & \mathcal{L}_{\mathrm{SS}} & \mathcal{L}_{\mathrm{SV}} \\ 0 & \mathcal{L}_{\mathrm{VS}} & \mathcal{L}_{\mathrm{VV}} \end{pmatrix} \cdot \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & \mathcal{P}_{\mathrm{VS}} & 0 \end{pmatrix} \cdot \begin{pmatrix} \vec{\psi}_{\mathrm{I}} \\ \vec{\psi}_{\mathrm{S}} \\ 0 \end{pmatrix} = \begin{pmatrix} \vec{s}_{\mathrm{I}} \\ \vec{s}_{\mathrm{S}} \\ 0 \end{pmatrix}. \tag{3.89}$$

Next, we explicitly carry out the matrix–matrix product and furthermore formally eliminate all degrees of freedom on the virtual domain Ω_V:

$$\begin{pmatrix} \mathcal{L}_\mathrm{II} & \mathcal{L}_\mathrm{IS} \\ \mathcal{L}_\mathrm{SI} & \mathcal{L}_\mathrm{SS} + \mathcal{L}_\mathrm{SV}\mathcal{P}_\mathrm{VS} \end{pmatrix} \cdot \begin{pmatrix} \vec{\psi}_\mathrm{I} \\ \vec{\psi}_\mathrm{S} \end{pmatrix} = \begin{pmatrix} \vec{s}_\mathrm{I} \\ \vec{s}_\mathrm{S} \end{pmatrix} . \tag{3.90}$$

This method has been (and still is) the standard approach for terminating wave guides in frequency–domain simulations. Within the accuracy of the finite Wannier expansion, this approach is exact and the performance is only limited by numerical imperfections such as round–off errors. It was validated by Daniel Hermann [30].

Reduced Basis Transformation Absorbers

The basis transformation absorbers presented in Sec. 3.5.4 work very well in frequency–domain problems but suffer from a practical imperfection. In Eqn. (3.88), *all* eigenmodes of the discretized wave guide contribute, in particular all evanescently decaying ones. Consequently, the full wave guide problem must be diagonalized, which scales cubically with the discretization.

However, if the device contains a sufficiently long piece of wave guide before the terminating stub, only propagating modes remain in this area and, consequently, only propagating modes must be continued by the propagator. Thus, we may approximate it by only retaining the parts of the eigenmode matrices $\mathcal{E}_\mathrm{WG}$ in Eqn. (3.88) that correspond to propagating modes and replace the rest of the matrix with a unit matrix. Furthermore, we must provide a “phase factor” for each erased evanescent mode. We use the number zero which means that these “modes” are not extended to the virtual wave guide area. The simplified propagator reads

$$\mathcal{P}_\mathrm{VS} \approx \underbrace{\begin{pmatrix} (\vec{\psi}_1)^\mathrm{T} \\ \vdots \\ (\vec{\psi}_M)^\mathrm{T} \\ \delta_{ij} \end{pmatrix}}_{\mathcal{E}_\mathrm{WG,red}} \cdot \begin{pmatrix} \mathrm{e}^{\mathrm{i}k_1\Delta R} & 0 & \\ & \ddots & \\ 0 & \mathrm{e}^{\mathrm{i}k_M\Delta R} & \\ & & 0 \end{pmatrix} \cdot \begin{pmatrix} (\vec{\psi}_1)^\mathrm{T} \\ \vdots \\ (\vec{\psi}_M)^\mathrm{T} \\ \delta_{ij} \end{pmatrix}^{-1} , \tag{3.91}$$

where $\vec{\psi}_1 \dots \vec{\psi}_M$ and $k_1 \dots k_M$ are the Wannier coefficients of all outgoing propagating eigenmodes and their propagating constants, respectively.

The unit matrix entries δ_{ij} in Eqn. (3.91) serve the single purpose of making the “reduced eigenmode matrix” $\mathcal{E}_\mathrm{WG,red}$ (which of course is no longer unitary) invertible. Depending on the profiles of the propagating eigenmodes of the wave guide, the condition of this matrix can become very poor which results in increased round–off errors and a degradation of the absorbing boundary. In such a case, the unit matrix part must be replaced with something more appropriate.

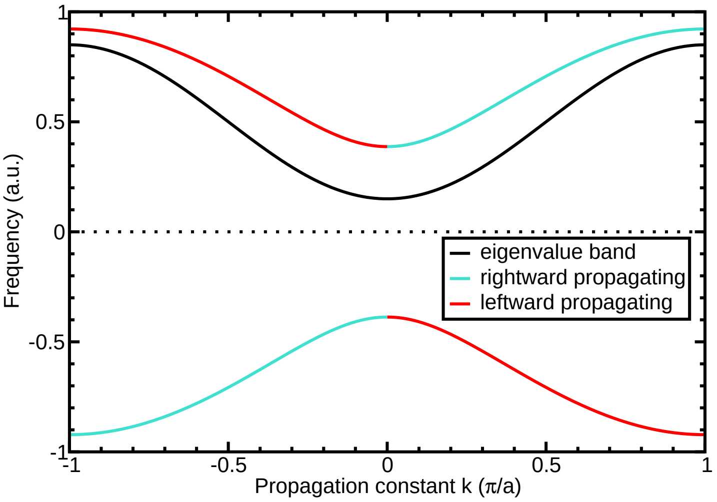

Figure 3.16.: *Eigenvalue band (black) and both branches of the dispersion relation of a single–band next–neighbor tight–binding wave guide. The red and cyan sections of the dispersion relation indicate leftward and rightward propagating modes, respectively. For an absorbing boundary condition based on a one–way wave equation one and only one of the red or the cyan frequencies must be reproduced. This requires slowly decaying long–range interaction terms.*

These simplified boundary conditions only require the propagating eigenmodes of the wave guide. In turn, however, they terminate only these modes and evanescent fields are reflected with erratic phases. Therefore, the simulated device should contain a sufficiently long piece of wave guide before the port. In this wave guide section, all field contributions except for the propagating modes decay exponentially, which drastically reduces the error introduced by the approximation Eqn. (3.91). If used properly, the performance of this simplified method is comparable to the full basis transformation absorbers.

One–Way Equation Absorbers

In the one–dimensional finite difference time–domain method, absorbing boundaries can be constructed by prohibiting any reflected waves to travel inside a boundary region. This is done by replacing the wave equation in this region with a one–way wave equation which only supports wave propagation in one direction [9].

Naturally, the question arises whether it is possible to construct a semi–discrete wave equation in the Wannier basis that only supports the outwards propagating eigenmodes of the wave guide. Such a formulation would have the great advantage of being non–dispersive which would make it suitable both for frequency–domain and time–domain simulations. Consequently, it is convenient to derive the equation in the frequency–

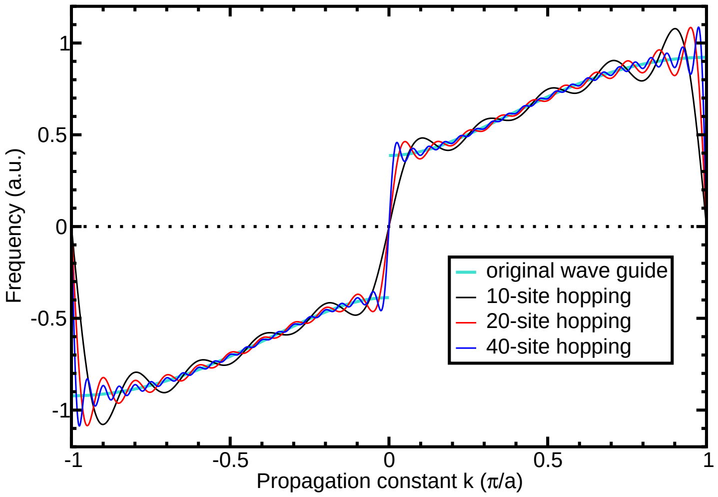

Figure 3.17.: *Dispersion relations of one–way equations that that were constructed to mimic the dispersion relation of the rightward propagating modes of a simple next neighbor tight–binding wave guide. Even if distant hopping over 40 lattice sites is taken into account, the group velocity mismatch is dramatic, in particular around the band edges.*

domain; the transformation to time–domain is straightforward. Unfortunately, one–way absorbing boundaries for periodic wave guides do not provide fundamentally better performance than the conceptually very simple adiabatic absorbers and are a dead–end. In the remainder of this section, I will assume that the outgoing waves travel rightward and the reflected waves travel leftward.

The full, lossless wave equation in semi–discrete form can be written as:

$$\partial_t^2 c_{n\mathbf{R}} = \sum_{n'\mathbf{R}'} \mathcal{L}^{(\text{full})}_{nn',\mathbf{R}\mathbf{R}'} c_{n'\mathbf{R}'}, \tag{3.92}$$

where (in a magnetic field formulation) $\mathcal{L}$ is the system matrix that results from multiplying Eqn. (3.68) with the inverse of $\mathcal{A}$[17]:

$$\mathcal{L} = \mathcal{A}^{-1}(\mathcal{C} + \mathcal{D}). \tag{3.93}$$

For the sake of simplicity, let us consider a single band, next–neighbor tight–binding model for the wave guide. This is a justified simplification because it models a coupled resonator optical wave guide (CROW), which has turned out to be difficult to terminate using an adiabatic absorber. Furthermore, the fundamental problems that will come up exist irrespective of whether more bands or long–range coupling terms are added to the original wave guide.

---

[17] The generalization to the electric field expansion is trivial: $\mathcal{L} = (\mathcal{C} + \mathcal{D})^{-1}\mathcal{A}$.

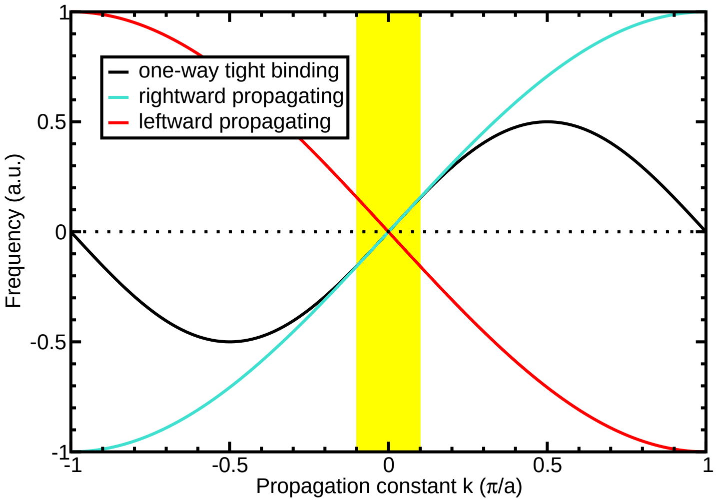

Figure 3.18.: *Both branches of the dispersion relation of a single–band next–neighbor tight–binding wave guide whose on–site term was chosen such that the positive branch and the negative one touch. In this case, the dispersion of all modes traveling in one direction is smooth around $k = 0$ and can be approximated for small wave vectors (yellow shaded area) by a simple sine.*

We restrict Eqn. (3.92) to only one band and transform it to frequency–domain:

$$\sum_{\mathbf{R}'} \mathcal{L}^{(\text{full})}_{\mathbf{R}-\mathbf{R}'} c_{\mathbf{R}'} = -\omega^2 c_{\mathbf{R}} = -\lambda c_{\mathbf{R}}. \tag{3.94}$$

A prototypical dispersion relation of this model as well as the eigenvalue $\lambda = \omega^2$ as a function of k are depicted in Fig. 3.16. In this single–band approximation, the system matrix $\mathcal{L}^{(\text{full})}_{\mathbf{R}-\mathbf{R}'}$ becomes a scalar, discrete function of $\Delta\mathbf{R} = \mathbf{R} - \mathbf{R}'$ and the eigenvalue band being its Fourier transform is a cosine.

As there are generally two solutions ω to the equation $\omega^2 = \lambda$, the dispersion relation consists of two branches, each of which is periodic in k–space[18]. Half of each branch corresponds to rightward propagating modes and the other half to leftward propagating ones. Consequently, a one–way equation would have to provide a dispersion relation that consists of the left half of the negative branch and the right half of the positive one (plotted in cyan in Fig. 3.16).

---

[18] These two branches are required to construct real–valued propagating waves by superposing modes with opposite frequency and wave vector:

$$\cos(kR - \omega t) = \frac{\exp\left(\mathrm{i}kR - \mathrm{i}\omega t\right) + \exp\left(-\mathrm{i}kR + \mathrm{i}\omega t\right)}{2}.$$

A one–way wave equation only contains first order temporal derivatives[19]. Consequently, we choose the ansatz:

$$\partial_t c_{\mathbf{R}} = \sum_{\mathbf{R}'} \mathcal{L}^{(\text{1way})}_{\mathbf{R}-\mathbf{R}'} c_{\mathbf{R}'}, \tag{3.95}$$

where $\mathcal{L}^{(\text{1way})}$ again is a scalar, discrete function of $\Delta\mathbf{R}$. In contrast to Eqn. (3.92), the dispersion relation of this wave equation is the Fourier transform itself of the system "matrix" rather than the square root. This means that the one–way system "matrix" has to be the inverse Fourier transform of the cyan curve in Fig. 3.16. As this curve is discontinuous, the hopping terms of $\mathcal{L}^{(\text{1way})}$ decay with first order and the reconstructed dispersion relation of any finite approximation will either contain oscillations (known as the Gibbs phenomenon) in both slow light regimes or will be smeared out if the long–range interaction is artificially damped in order to suppress the oscillations.

As a consequence, the error introduced by this type of absorbers is reciprocal to the absorber length which is not superior to adiabatic absorbers. Furthermore, one–way absorbers for wave guides can be expected to massively distort the slow–light regime. How severe this problem is can be seen in Fig. 3.17, where I plotted the dispersion relations for one–way equations with the long–range interaction restricted to up to 40 lattice sites. This is already more than what is required to obtain a fairly usable adiabatic absorber for a rather broad–band wave guide. For these reasons I discarded this ansatz.

For completeness, let us consider the special case that the dispersion relation of the full wave guide reaches zero at $k = 0$. In this case, the dispersion relation of the full wave guide consists of two sine branches each of which is periodic on twice the Brillouin zone (cf. Fig. 3.18). An appropriately chosen strict next–neighbor coupling in Eqn. (3.95) yields a sine dispersion relation (black curve in Fig. 3.18) which is periodic on the Brillouin zone. Although this sine massively deviates from the full dispersion relation close to the boundaries of the Brillouin zone, it exhibits the same linear behavior around $k = 0$. Thus, it can act as an acceptable absorbing boundary condition provided the wave length of the propagating wave is much larger than the unit cell of the lattice. This would be the situation in a typical finite–difference time–domain [9] calculation.

Time–Domain Convolution Absorbers

The elegance and excellent performance of the basis–transformation absorbers in frequency–domain suggest to extend this approach to time–domain. However, there is the fundamental problem that the propagator $\mathcal{P}_{\text{VS}}$ strongly depends on frequency. Consequently,

---

[19] As the derivative ∂_t^n is translated to $(-\mathrm{i}\omega)^n$ under a Fourier transform, the differential equation order equals the number of branches in the dispersion relation; the eigenfrequencies are the roots of the equation $\omega^n = \lambda$, where λ is an appropriately defined eigenvalue. Thus, we could also aim for a third order differential equation which would provide two extremely lossy bands in addition to the desired rightward propagating one. Unfortunately, this does not avoid the problems that also occur for the first order ansatz.

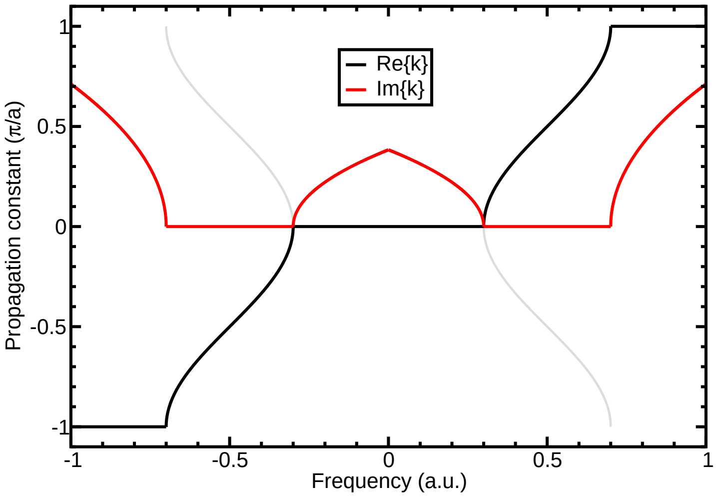

Figure 3.19.: *Qualitative sketch of the propagation constant function that must be Fourier transformed to yield an appropriate convolution kernel. Within the bandwidth of the wave guide, the propagation constant is given by the dispersion relation of the rightward propagating wave guide modes (cf. Fig. 3.16; for comparison, the leftward propagating branches were added to this plot in light gray). In between, the propagation constant acquires an imaginary part that causes an exponential decay of the evanescent modes. The imaginary part in this plot is not properly scaled with respect to the real part.*

the simple frequency–domain product

$$\mathcal{P}_{\mathrm{VS}}(\omega)\vec{\psi}_{\mathrm{S}}(\omega) \tag{3.96}$$

in Eqn. (3.90) translates to a convolution in time–domain:

$$\int_0^t \mathrm{d}t' \mathcal{P}_{\mathrm{VS}}(t-t')\vec{\psi}_{\mathrm{S}}(t'), \tag{3.97}$$

where $\mathcal{P}_{\mathrm{VS}}(\tau)$ is the time–domain wave guide propagator:

$$\mathcal{P}_{\mathrm{VS}}(\tau) = \int_{-\infty}^{\infty} \mathrm{d}\omega \, \exp\left(\mathrm{i}\omega\tau\right) \mathcal{P}_{\mathrm{VS}}(\omega). \tag{3.98}$$

For a single–band tight–binding wave guide, $\mathcal{P}_{\mathrm{VS}}(\omega)$ qualitatively resembles the function sketched in Fig. 3.19. The first derivative of $k(\omega)$ is not only discontinuous but not even bound at the discontinuity. Consequently, the time–domain propagator $\mathcal{P}_{\mathrm{VS}}(\tau)$ decays very slowly with respect to τ. This is a direct consequence of the existence of slow light modes. Consequently, modifying the lower boundary in the integral Eqn. (3.97) results in a relatively large error unless the upper bound t is very large.

On the other hand, the convolution must be performed over a finite and fixed time interval for practical reasons: First, the most straightforward way to calculate $\mathcal{P}_{\mathrm{VS}}(\tau)$ is from the frequency–domain propagator via Eqn. (3.98). However, the finite frequency spacing of the numerically determined $\mathcal{P}_{\mathrm{VS}}(\omega)$ causes an artificial periodicity with respect to time. Furthermore, keeping the fixed lower integration boundary in Eqn. (3.97) would result in the computational time to scale quadratically with respect to the simulated time, which is unacceptable in practice.

I performed tests with as many as 10000 intermediate steps in Eqn. (3.97) without obtaining even remotely acceptable wave guide termination. This requires 9999 additional sets of expansion coefficients that describe the previous states and as many values for $\mathcal{P}_{\mathrm{VS}}(\tau)$, each of which is an individual matrix. Consequently, the memory and time consumption of these tests were completely unacceptable.

Outlook: Non–Reciprocal Wave Guides

A promising ansatz for efficient absorbing boundaries occurred to me during a discussion with one of our diploma students. The basic idea is somewhat similar to the one–way equation boundary conditions. Excluding the leftward propagating modes creates discontinuities in the dispersion relation that lead to poor performance. However, we may construct a second–order differential equation whose dispersion relation is asymmetric with respect to the mapping $k \mapsto -k$. More precisely, we may try to construct a wave guide which exactly mimics the original wave guide for rightward propagating modes but whose leftward propagating modes are subject to strong losses due to an imaginary part of the frequency $\omega(k)$. This way, the light travels through the absorber, is reflected at the far side and attenuated as it travels back towards the simulation domain.

This ansatz avoids discontinuities in the dispersion relation, which provides the prospect of at least second–order convergence with respect to the absorber length. The imaginary part of ω causes the reflected fields to decay with a certain *rate*. We obtain an equal attenuation for all reflected modes if this rate is reciprocal to the time each mode spends inside the absorber. If we consider a narrow–spectrum pulse around one mode, this time is reciprocal to the group velocity of the center mode. Thus, in order to guarantee a minimum attenuation for all modes, there must exist a global constant γ such that

$$\exists \gamma \in \mathbb{R}^+ : \quad \mathrm{Im}\{\omega(k)\} \geq \gamma \partial_k \mathrm{Re}\{\omega(k)\} \quad \text{for leftward propagating modes.} \tag{3.99}$$

We may expect optimal performance if the dispersion relation fulfills above relation as an equation rather than an inequality.

4 Chapter 4.

Transmission Through Opal Films

As a first practical project I would like to present an investigation of the transmission of light through thin opal films. In particular, I discuss the change of the polarization state of light impinging under oblique incidence upon transmission. This work was done in close collaboration with Prof. Sergei G. Romanov at the University of Erlangen, who performed the experiments, and Jens Küchenmeister from our group, who contributed calculations based on the Fourier Modal Method (FMM) [11, 12]. For all band structure calculations, I used the software package MPB [16].

The different numerical methods shed light on very different aspects of the system. The FMM simulates the physical system with all its finite–size and surface[1] effects. The problem of convergence left aside, the spectra obtained with this method would be exactly the same as the measured ones if the experiments could be performed with an ideal sample (no surface roughness, no crystal disorder). Thus, it offers quantitative predictions but it is not optimal to gain insight in the underlying physics. The band structure on the other hand provides the propagating modes of the infinite bulk photonic crystal (PhC), their frequencies and their wave vectors. With this knowledge it is possible to decompose the process of refraction at a PhC interface into several mechanisms and to analyze them independently. Knowing the band structure, it is possible to tell to which propagating modes a plane wave can couple and it is possible to investigate those modes independently. However, it is very difficult to quantitatively predict transmission or reflection coefficients. Oversimplifying the picture, one could say that the FMM reproduces the outcome of an experiment and the band structure can provide an explanation.

After a short introduction, I present the experimental setting and show some measured spectra together with ab–initio calculations performed with the FMM. These plots exhibit rather unexpected polarization effects and mysteriously changing transmission into different Bragg orders. In the next section, I discuss refraction at a PhC interface and how it is connected to the band structure. The remainder of this chapter deals with

---

[1] This does not refer to any roughness of the spheres but to the grating–like nature and the specific structure of the PhC interface.

the polarization effect in which I compare my analysis to the FMM calculations. I do not directly compare to experimental data because they exhibit specific strong deviations due to fabricational imperfections, which I address in the last section of this chapter.

4.1. Introduction

Thin films of artificial opals are a class of photonic crystal structures that consist of dielectric spheres arranged in a face–centered cubic (fcc) lattice. Although these systems do not exhibit a photonic band gap, they provide a platform for investigating fundamental properties of photonic crystals in general.

Transmittance minima are the most striking features of such films and are often related to photonic stop bands [52, 53]. More information can be gained by additionally measuring the polarization state of the reflected and transmitted light [54]. For instance, the PhC film can alter the polarization state of the incident radiation. This type of optical anisotropy in photonic crystals has been known for over a decade in macroporous silicon [55] and was recently found in strained opal films [56]. In the former case, the system geometry is extremely anisotropic simply because the PhC features only a two–dimensional periodicity. In the latter case, the cubic symmetry of the opal is broken by strain. In both cases, the polarization transfer is (at least also) visible in the long wave length limit and, thus, can be explained by conventional form birefringence [57].

However, nontrivial polarization dynamics in a system with the full cubic symmetry was noticed in simulations some years [58] ago but not further investigated. This was one of the topics of our collaboration with Prof. Romanov and the subject of this chapter.

4.2. Experimental Findings and FMM–Calculations

The system investigated by Prof. Romanov is a thin opal film. Artificial opals are photonic crystals that consist of dielectric spheres in a face–centered cubic lattice. Relatively large samples can be fabricated by self–assembly of spheres dispersed in an appropriate solvent [54]. This cost–efficient low–tech[2] fabrication process together with the excellent quality of the samples made this class of photonic crystals a popular platform for fundamental research of PhCs despite the lack of a complete photonic band gap.

In the experiment, the opal film made of polymethyl–metacrylate ($\epsilon \approx 2.22$) spheres (thickness typically around 25 layers of spheres) is illuminated with polarized light under oblique incidence (cf. Fig. 4.1). The transmitted and the reflected light is then decomposed into its respective Bragg orders and the frequency spectrum for both polarization states (either s–pol or p–pol) of each order is analyzed. By recording such spectra for varying angles of incidence, one obtains polarization–resolved reflectance and transmittance maps as depicted in Fig. 4.2. The experimental and calculated bar–polarized (s–pol into s–pol) transmittance plots are in excellent agreement expect for a factor

---

[2] as compared e. g. to electron beam lithography

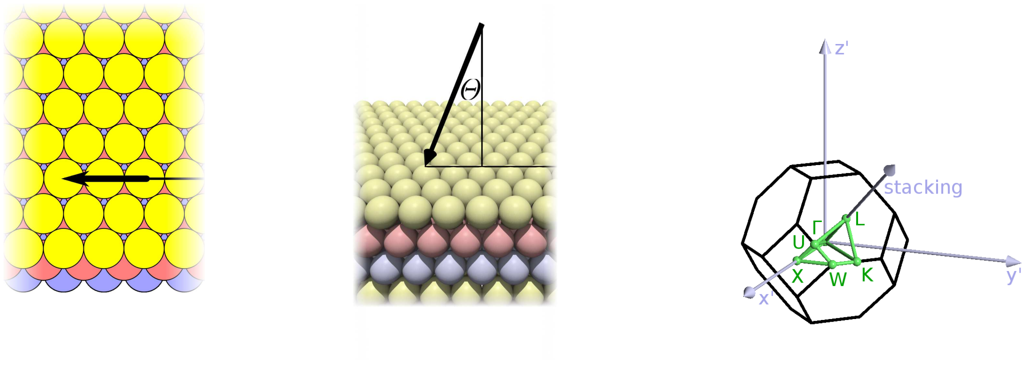

(a) top view (b) isometric view (c) first Brillouin zone

Figure 4.1.: *Simplified sketches of the experimental setup and first Brillouin zone of the underlying fcc–lattice. The opal film consists of layers of spheres stacked in the Γ–L direction of a fcc–lattice. The three types of layers are indicated by different colors. The wave vector of the incident light wave is indicated by a bold black arrow. In these sketches, the direction of incidence lies in the Γ–L–W plane. The two left pictures were taken from our submitted manuscript Ref. [59].*

of two, which Prof. Romanov attributes to scattering in the film due to fabricational imperfections.

The cross–polarized (s–pol into p–pol) spectra, however, differ by a factor of about six while the qualitative agreement is similarly good as in the bar–polarized case. As both pairs of spectra are measured in transmission and there is no real reason why scattering should be three times stronger for cross–polarization, some other effect must reduce the signal without qualitatively modifying the spectrum. For this, I will provide an explanation in Sec. 4.4.4.

4.3. Bloch Wave Picture of Refraction at PhC Interfaces

The polarization conversion effect appears in transmission. Consequently, we must understand light propagation in air, inside the PhC and in the glass substrate. The propagating modes inside the isotropic homogeneous media and inside the PhC are plane waves and the Bloch modes, respectively. The latter ones can be determined by a band structure calculation (cf. Chap. 2). However, which Bloch modes are relevant for this problem depends on the wave length and the propagation direction of the incident plane wave and on the shape and orientation of the PhC surface. In short, we must solve the refraction problem at a PhC interface.

Refraction at a planar interface of a homogeneous material is governed by four conditions:

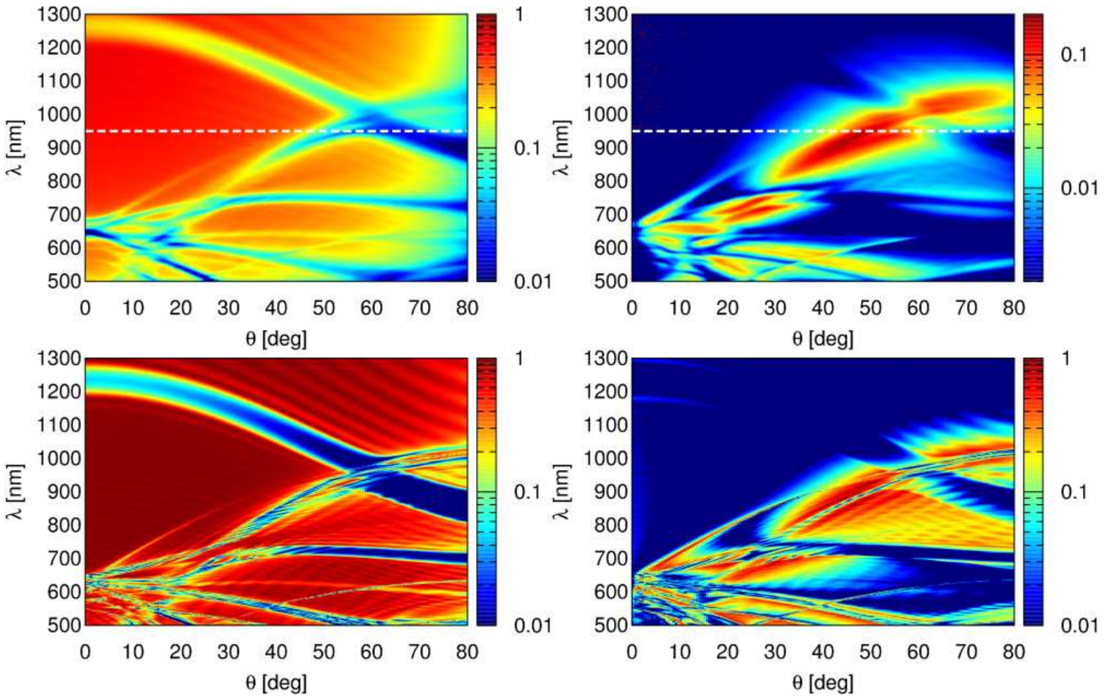

Figure 4.2.: *Experimental (top row) and calculated (bottom row) polarization–resolved transmittance maps for the zeroth Bragg order for s–polarized radiation incident in the* Γ*–*L*–*W *plane. The maps depict the relative power (color–coded) that is transmitted into the same polarization (s–pol into s–pol, bar-polarized transmittance, left column) and into the perpendicular polarization (s–pol into p–pol, cross-polarized transmittance, right column). This plot was taken from our submitted manuscript Ref. [59].*

1. Conservation of frequency: Requires the incident, transmitted and reflected waves have the same frequency.
2. Conservation of in–plane momentum: Defines the wave vectors of the transmitted and reflected beams.
3. The continuity conditions of the EM-fields at a material interface: Defines the ratio between transmitted and reflected radiation.
4. Conservation of energy: Requires that the total transmittance and total reflectance sum up to one (unless energy is dissipated at the interface).

The fourth point is very fundamental, of course, and conservation of energy is of utmost importance in any quantitative calculation. The band structure analysis presented in this chapter only provides qualitative insight in the dominant polarization conversion process

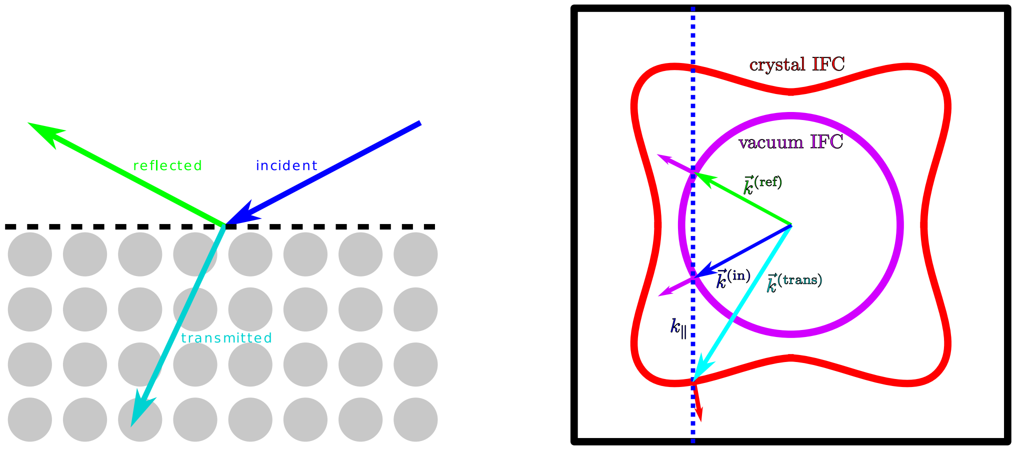

(a) refraction in real space (b) refraction in momentum space

Figure 4.3.: *Sketch of the refraction process at a photonic crystal interface (black dotted line) (a) in realspace and (b) in reciprocal space: The incident wave vector (blue arrows) is decomposed into the parts parallel and perpendicular to the interface (decomposition not shown). Possible* **k***–vectors in each material (vacuum and the PhC) can be found by intersecting the conserved* **k***–vector component (blue dotted line in the right panel) and the respective isofrequency contour (IFC, depicted in red for the PhC and in violet for vacuum). Furthermore, wave vectors for outgoing waves must correspond to an energy velocity (i. e. the isofrequency normal at the intersection; depicted by small arrows with the same color as the corresponding IFC) pointing into the respective material. Here, this results in one reflected wave (green) and a transmitted one (cyan).*

and does not aim to yield absolute transmission coefficients. Thus, we may ignore energy conservation. In the case of a PhC, the interface has a discrete translational symmetry which modifies two of the conditions stated above:

2a. The in–plane part of the wave vector is conserved only up to a reciprocal lattice vector. This leads to the occurrence of Bragg reflections.

3a. The continuity conditions have to be integrated over a unit cell of the periodic interface.

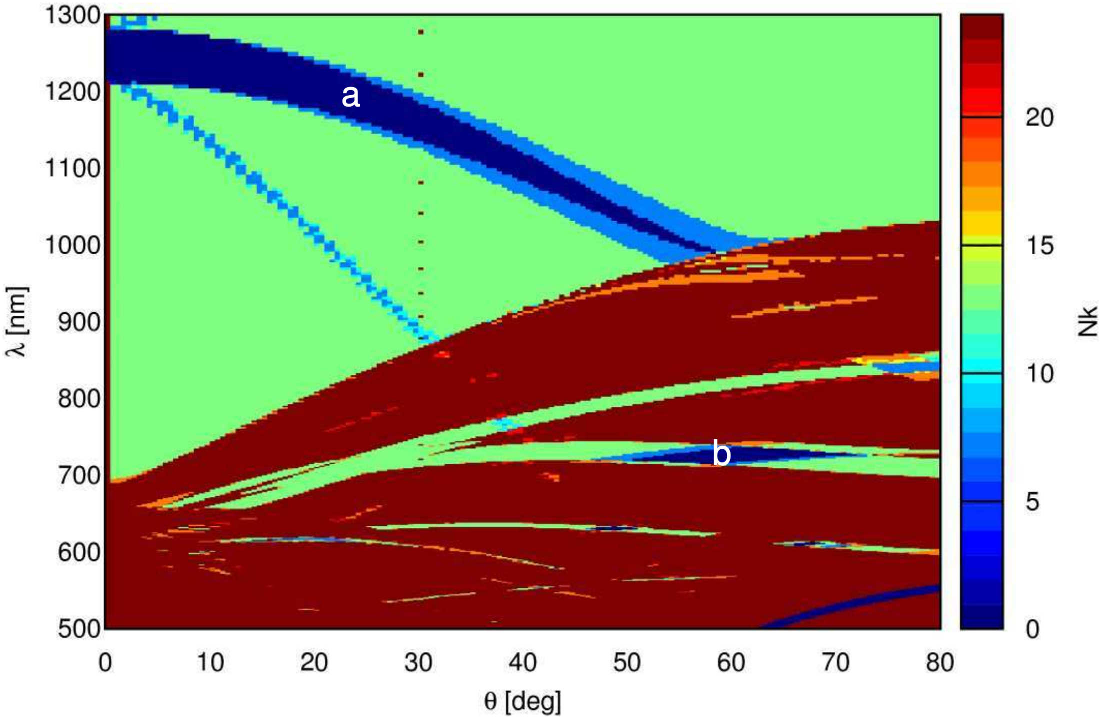

Figure 4.4.: *Color–coded map of the number of wave vectors allowed by conservation of momentum for various wave lengths and angles of incidence. One should keep in mind that due to the highly symmetric choice of the interface normal many reciprocal lattice vectors have the same* $\mathbf{G}_{\|}$ *and that thus the coupling numbers are artificially increased. Furthermore, forward and backward propagating modes both are counted. As a result, for wave length above the fundamental stop band, this plots shows twelve reachable wave vectors (light green) although there are only two unique propagating modes. The color axis is artificially saturated at 25. As a consequence, this plot only shows parameter ranges with no (blue), few (green) and many (red) reachable wave vectors and hides the rather complex structure (which is rather confusing than informative) at small wave lengths. For each point it is still to be determined whether the Bloch modes match the incident wave with respect to symmetry.*

4.3.1. Wave Vector Decomposition and Conservation

Let us consider a PhC interface with normal vector $\hat{n}$ and a plane wave in vacuum incident at some direction $\hat{e}_k$ with the wave vector

$$\mathbf{k}^{(\mathrm{in})} = \frac{\omega_0}{c}\hat{e}_k. \tag{4.1}$$

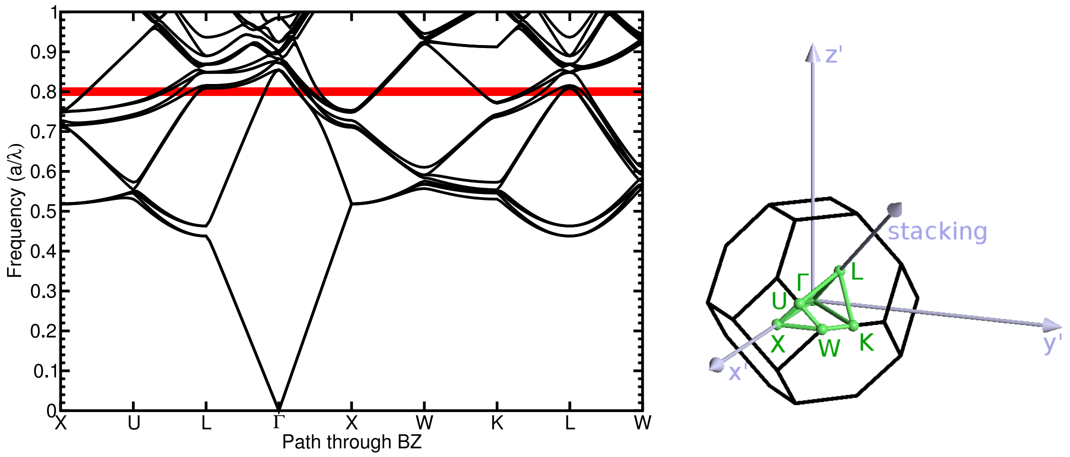

Figure 4.5.: *Photonic band structure for the opal system (PMMA spheres in a fcc–lattice) calculated using my implementation of the plane wave expansion at 7500 plane waves. The red shaded area corresponds to vacuum wave lengths around* 700 *nm where an additional transmission gap at great incident angles is visible in Fig. 4.2. In this band structure plot, there are no features visible which could explain the suppressed transmission in just this frequency window.*

In the interface's normal direction, translational symmetry is completely broken and, thus, the normal wave vector component

$$\mathbf{k}_\perp^{(\mathrm{in})} = \hat{n}\mathbf{k} \cdot \hat{n} \tag{4.2}$$

is not conserved. In the plane, however, there is a discrete translational symmetry which ensures that the in–plane part of the wave vector

$$\mathbf{k}_\parallel^{(\mathrm{in})} = \mathbf{k} - \mathbf{k}_\perp \tag{4.3}$$

is conserved up to the parallel component of a reciprocal lattice vector $\mathbf{G}_\parallel$ of the crystal.

Thus, conservation of momentum imposes the restriction that the incident plane wave can only couple to those Bloch modes that fulfill:

$$\begin{aligned} \omega_0 &= \omega^{(\mathrm{BF})}_{n\mathbf{k}^{(\mathrm{BF})}}, & (4.4)\\ \mathbf{k}_\parallel^{(\mathrm{in})} &= \mathbf{k}_\parallel^{(\mathrm{BF})} + \mathbf{G}_\parallel, & (4.5)\end{aligned}$$

where $\omega^{(\mathrm{BF})}_{n\mathbf{k}^{(\mathrm{BF})}}$ is the eigenfrequency of the Bloch function corresponding to band n and wave vector $\mathbf{k}^{(\mathrm{BF})}$. This is equivalent to intersecting the isofrequency contour (i. e. the set of all wave vectors that correspond to the same frequency) of the PhC with all lines in $\mathbf{k}$–space that are defined by

$$\mathbf{k}^{(\mathrm{BF})} - \hat{n}\mathbf{k}^{(\mathrm{BF})}\hat{n} = \mathbf{k}_\parallel^{(\mathrm{in})} - \mathbf{G}_\parallel = \mathrm{const.} \tag{4.6}$$

This is also illustrated in Fig. 4.3. For small angles of incidence and large wave lengths, fold–back by a reciprocal lattice vector is irrelevant because in this case, there are no intersection points with the isofrequency surface for $\mathbf{G}_{\|} \neq \mathbf{0}$.

This very basic mechanism depends only on the normal vector of the interface and, thus, is independent of the actual structure of the interface as long as it respects the PhCs periodicity.

In Fig. 4.4, I plotted the number of PhC wave vectors that fulfill Eqn. (4.6) for an interface normal to the Γ–L–direction and for incident waves in the Γ–W–plane. To this end, I performed a pseudo–on–shell calculation by linearly interpolating between the results of conventional band structure calculations as described in Sec. 2.3.3. The coupling gap marked with "a" corresponds the fundamental stop band between the second and the third photonic band and can thus be predicted by simple inspection of a conventional band structure plot. The coupling gap labeled "b" however corresponds to a spectral range without any stop band. There, the band structure (cf. Fig. 4.5) looks completely innocent with a continuum of states. There is not qualitative difference from the wave lengths above and below. The incident wave cannot couple because of the choice of the interface normal and the specific shape of the isofrequency contour at these frequencies. This means that the type "a" gap is fundamentally present due to a lack of states in the bulk crystal whereas the type "b" gap occurs due to the interface orientation and can (in principle) be removed by cutting the crystal along a different direction. However, it cannot be removed by modifying the interface's structure for a given crystal lattice and a given interface normal.

If we compare Fig. 4.4 with Fig. 4.2, we find that strong cross–polarized transmission occurs in a parameter regime where coupling to additional Bloch modes is possible. This already suggests that it could be caused by interference effects.

For each point of the plot Fig. 4.4, it is still to be determined whether the Bloch modes match the incident wave with respect to symmetry (cf. Sec. 4.3.2).

4.3.2. Bloch Wave Symmetries

The above considerations depend on the interface normal vector only and predict whether coupling to that mode is prohibited by conservation of momentum and frequency. This result can be further refined by mode–matching the plane wave and the Bloch functions. The coupling amplitude itself is given by the overlap between the Bloch mode's field distribution and the incident wave:

$$t'_{n\mathbf{k}} = \int_{\mathrm{IUC}} \mathrm{d}^2 r \left(\mathbf{H}^{(\mathrm{BF})}_{n\mathbf{k}^{(\mathrm{BF})}}(\mathbf{r}) \right)^* \left(\mathbf{H}^{(\mathrm{in})}(\mathbf{r}) \right), \tag{4.7}$$

where IUC refers to the unit cell of the interface and the incident field is given by

$$\mathbf{H}^{(\mathrm{in})}(\mathbf{r}) = \mathbf{H}^{(\mathrm{in})}_0 \exp\left(-\mathrm{i}\mathbf{k}^{(\mathrm{in})}\mathbf{r} \right). \tag{4.8}$$

We assumed that the permeability $\mu = 1$ in all materials. We note that for the zeroth Bragg order, the Bloch phase in the plane takes the same values as the incident wave

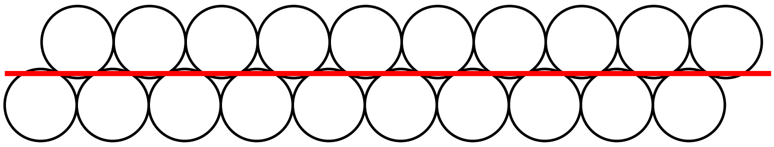

Figure 4.6.: *Sketch of the termination plane used for mode matching the Bloch modes and the incident wave. The plane is in Γ–L–direction such that it is as far away from any sphere center as possible. Although it cuts all spheres of both adjacent layers, it is the best compromise, in my opinion.*

.

due to conservation of momentum. Thus, the quotient between those phases is one:

$$\exp\left(\mathrm{i}\mathbf{k}^{(\mathrm{BF})}\mathbf{r} - \mathrm{i}\mathbf{k}^{(\mathrm{in})}\mathbf{r}\right) = 1 \quad \forall \mathbf{r} \in \mathrm{IUC}. \tag{4.9}$$

For higher Bragg orders, the reciprocal lattice vector that was used to connect $\mathbf{k}^{(\mathrm{in})}$ and $\mathbf{k}^{(\mathrm{BF})}$ remains

$$\exp\left(\mathrm{i}\mathbf{k}^{(\mathrm{BF})}\mathbf{r} - \mathrm{i}\mathbf{k}^{(\mathrm{in})}\mathbf{r}\right) = \exp\left(-\mathrm{i}\mathbf{G}_{\|}\mathbf{r}\right) \quad \forall \mathbf{r} \in \mathrm{IUC}. \tag{4.10}$$

Thus, we may simplify Eqn. (4.7):

$$t'_{n\mathbf{k}} = \mathbf{H}_0^{(\mathrm{in})} \int_{\mathrm{IUC}} \mathrm{d}^2 r \left(\mathbf{h}^{(\mathrm{BF})}_{n\mathbf{k}^{(\mathrm{BF})}}(\mathbf{r})\right)^* \exp\left(\mathrm{i}\mathbf{G}_{\|}\mathbf{r}\right) \tag{4.11}$$

where $\mathbf{h}^{(\mathrm{BF})}_{n\mathbf{k}^{(\mathrm{BF})}}(\mathbf{r}) = \exp\left(\mathrm{i}\mathbf{k}^{(\mathrm{BF})}\mathbf{r}\right)\mathbf{H}^{(\mathrm{BF})}_{n\mathbf{k}^{(\mathrm{BF})}}(\mathbf{r})$ is the lattice–periodic part of the Bloch function. In a nutshell, we simply mode–match the lattice–periodic parts of both waves rather than the waves themselves. The advantage is that the lattice–periodic part of a plane wave is a constant and thus the coupling can be estimated by simply looking at the Bloch modes' periodic field pattern.

The expression above is the exact coupling coefficient to the Bloch mode provided that the crystal is terminated by a straight cut plane. Although this is not the case for opal films, we may hope that many symmetry properties of the Bloch modes are robust and do not change dramatically if the interface structure is modified. For our specific problem, we chose a cut plane normal to the Γ–L–direction such that the plane is as far away as possible from the sphere centers of two adjacent layers (cf. Fig. 4.6). This plane is an acceptable approximation to the experimental setup.

To illustrate the mode matching problem, I plotted the lattice–periodic part of a magnetic Bloch mode on the IUC of said plane in Fig. 4.7. The phase plots reveal that the z–component of this mode cannot couple to a zeroth Bragg order wave because the periodic phase change makes the coupling integral vanish. However, the phase plots of the x– and y–components are flat and, thus, these components can couple to such an incident wave. Furthermore, the phase difference between x– and y–component is approximately π, which means that the in–plane component of this mode at the interface resembles a linearly polarized wave. Please note that this does not tell anything about the “polarization state” of this mode anywhere else in the unit cell. Moreover, the wave *is not* linearly polarized in the cut plane because the z–component exhibits a complete phase cycle. It just couples to zeroth Bragg order waves as if it was linearly polarized.

The phase accumulated throughout the unit cell is a very robust property of the Bloch mode so we can deduce that coupling to a mode with an inappropriate phase distribution is highly suppressed at the experimentally relevant interface as well.

4.3.3. Refraction and Band Structure in a Nutshell

Band structure calculations can predict the transmission through a PhC on different levels of accuracy depending on the amount of information provided about the actual system. Given only the crystal structure, total band gaps give a first hint. If the direction of propagation inside the crystal is known, stop bands refine the picture. If the orientation of the interface is known, transmission can be predicted to vanish based on conservation of momentum even if there are no remarkable features in the band structure plot. Finally, mode matching can exclude further possible propagation channels based on symmetry considerations and it can approximate transmission coefficients.

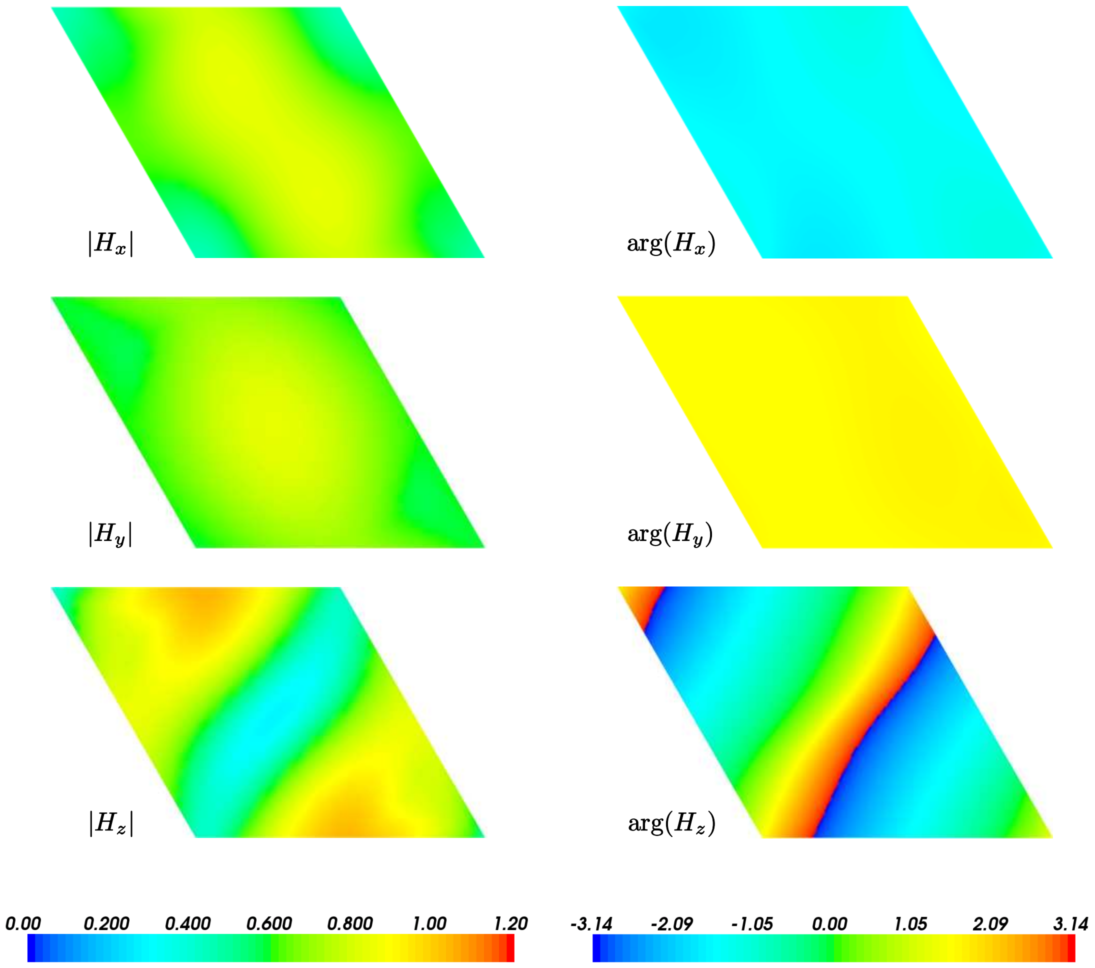

Figure 4.7.: *Cut parallel to the PhC–interface through the lattice–periodic part of the magnetic Bloch mode band 5 at* $\mathbf{k}$*–vector* $(-0.224, -0.448, 0.190)$ *in reciprocal lattice coordinates. Left column: Absolute value of each field vector component; right column: corresponding phase of each field vector component. From top to bottom:* H_x *(in the interface plane and in the plane of incidence),* H_y *(normal to the plane of incidence),* H_z *(normal to the interface plane). Based on momentum conservation at the interface, a wave (*$\lambda = 900\,nm$*) incident at an angle of* $47°$ *(zeroth Bragg order) could couple to this mode.* H_x *and* H_y *have flat phase patterns whereas* H_z *exhibits a complete phase change. Thus, a plane wave can couple to the in–plane field components of this mode. The phase difference between* H_x *and* H_y *is nearly* π *at each point which indicates that the in–plane part of this mode is linearly polarized.*

4.4. Effective Birefringence

Experiments and FMM simulations (Fig. 4.2) show that there exists a regime where the polarization state of light is altered upon transmission through the opal film. In this section, I describe the approach I used to understand this effect and present my findings that led to our manuscript Ref. [59].

In principle, the modification in the polarization state might be caused either by a surface effect, a bulk effect or a combination thereof. For closer investigation we restrict ourselves to the case of an incident plane wave with a wave length of about 900 nm and tilt angles between 42 and 52 degrees in the Γ–L–W-plane (cf. Fig. 4.1 for a sketch of the first Brillouin zone). In this range of parameters, higher Bragg orders do not couple which makes the analysis somewhat easier.

First, we employed the FMM to compute the polarization–resolved transmission coefficients over the stated range of parameters for various opal film thicknesses. As bulk effects are expected to feature a strong dependence on the film thickness and surface effects are not, these calculations should give a serious hint about the nature of the polarization transfer. In particular, bulk effects would lead to interference beat patterns for the cross–polarized (e. g. s–pol into p–pol) and the bar–polarized (e. g. s–pol into s–pol) transmittance, whereas surface effects would cause a modification that is roughly independent of the film thickness. In the following, I present my analysis in detail only for $\Theta = 47°$.

4.4.1. FMM–Results

First, we should note that the FMM–calculations include distracting features caused by the finite opal film thickness, most prominently Fabry–Pérot resonances which lead to fluctuations on top of the transmittance spectra. Furthermore, as the number of sphere layers in a film is a small integer, the simulation data points are given at a relatively coarse spacing. Therefore, it is impossible to derive sufficiently accurate quantities directly from the raw ab–initio results. Instead, we must fit an appropriate sinusoidal curve to the simulation results.

Transmittance Spectra

In Fig. 4.8, we show typical[3] bar–transmission (s–pol into s–pol) and cross–transmission (s–pol into p–pol) spectra. It is easily seen that both transmission spectra as functions of the film thickness include zero or at least become very small. This holds for all other

---

[3] Violating a widespread habit throughout science, in this context the word “typical” is not synonymous to “particularly beautiful”.

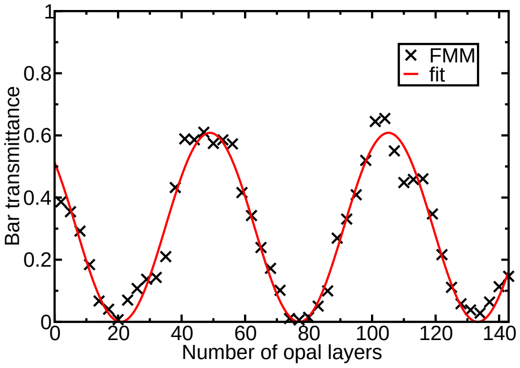

(a) s–pol into s–pol

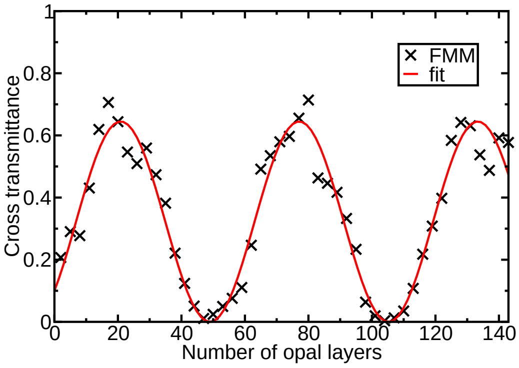

(b) s–pol into p–pol

Figure 4.8.: *Oscillations of the polarization–resolved transmittance with respect to the film thickness. Black crosses: Bar–transmittance (s–pol into s–pol, left panel) and cross–transmittance (s–pol into p–pol, right panel) computed by the FMM for an incident wave at a wave length of* 900 *nm and an incident angle of 47 degrees in the* Γ*–L–W–plane. Red line: Sinusoidal fit to the FMM–results. The ab–initio results deviate from the fit due to Fabry–Pérot resonances. FMM data are courtesy of Jens Küchenmeister.*

angles as well. Thus, we may choose a fit ansatz with three free parameters

$$T(d) = \alpha\Big(\sin(2\pi\beta d+\gamma)+1\Big), \tag{4.12}$$

$$\alpha \in [0,0.5], \tag{4.13}$$

$$\beta \in [0,\infty), \tag{4.14}$$

$$\gamma \in [0,2\pi), \tag{4.15}$$

where d denotes the film thickness. The least square fits of this function to the FMM–results are plotted in red in Fig. 4.8.

The oscillations in the bar–transmittance and the cross–transmittance plots have very similar periods and amplitudes but, to an excellent approximation, differ in phase by π. Furthermore, the oscillations have a very large period (around 50 layers of spheres), which makes it very unlikely that they are caused by Fabry–Pérot–like interference effects. Instead, we can safely assume that the polarization transfer is dominated by an interference beat between PhC bulk modes and that the surface does not play a major role.

Transmitted Polarization State

So far, it is only clear that the transmitted light is a superposition of s–pol and p–pol but the phase relation is yet to be determined. To this end, we computed the complex cross–

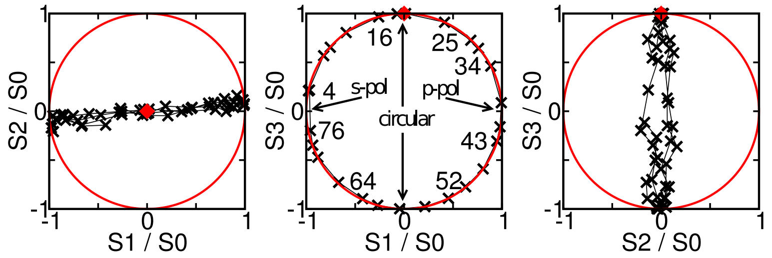

Figure 4.9.: *Stokes parameters (S_0 normalized to one) of the transmitted polarization states for various film thicknesses (black crosses) at* 900 nm *wave length and an incident angle of 45 degrees. I plotted three views: Parallel to S_3 (left plot), parallel to S_2 (center plot) and parallel to S_1 (right plot). The "north pole" of the sphere is indicated by a red diamond. The states lie on the surface of the Poincaré–sphere close to a circle through its poles. This is the typical fingerprint of a birefringent material. The number labels in the central panel indicate the film thickness (in units of sphere layers) for which the adjacent data point was calculated.*

and bar–transmission coefficients. The Stokes parameters of the resulting polarization states in relation to the Poincaré–sphere are shown in Fig. 4.9. The polarization states lie on the sphere's surface and are close to a circle through both poles. This means that starting with a very thin film, the transmitted light (for instance) is s–polarized, becomes elliptical as the film thickness is increased, reaches near–circular polarization, then approaches p–polarization and for even thicker films transforms back again. This is a typical behavior for a birefringent material. The complete presence of linear anisotropy rather than chiral effects might be slightly surprising because the material itself is chiral due to the stacking order "ABC..." but has full cubic symmetry. Hence, linear (form) birefringence naively would not be expected [57]. This is only possible because we are far away from the effective medium limit.

4.4.2. Band Structure Cuts

Now, we can interpret the FMM–results based on band structure considerations. Usually, band structures are plotted on the outline of the irreducible wedge of the Brillouin zone. This path, however, is inappropriate for our question as it does not contain the **k**–vectors of the Bloch modes to which coupling is possible. Instead, we plot the band structure on a path that is defined by the condition $\mathbf{k}_{\|} = \text{const}$, i. e. on the blue dotted line in Fig. 4.3. I would like to stress that for a given angle of incidence, the cut is valid for one frequency, only. This frequency is indicated in the band structure cut plots as a solid red line. Such plots for 900 nm and 47° or 51° are depicted in Fig. 4.10. Every intersection of the horizontal red line with a photonic band (solid black lines) defines one

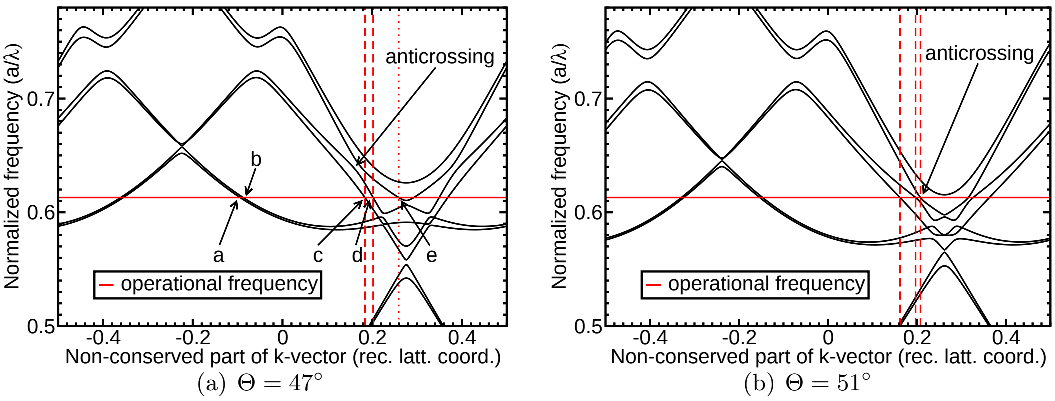

(a) $\Theta = 47°$ (b) $\Theta = 51°$

Figure 4.10.: *Band structure cuts along the line of* $\mathbf{k}_{\|}$ *defined by the zeroth order of a wave at incident angles of 47 degrees (left panel) and 51 degrees (right panel). In the left panel, there are five relevant intersection points of the operational frequency and the band structure, labeled with 'a' through 'e'. The other five intersection points are identical to the former ones up to a symmetry operation and describe modes traveling in the opposite direction. An important feature in this band structure plot is the band anticrossing. It moves to lower frequencies as the angle of incidence is increased and reaches the operational frequency for 51 degrees (right panel).*

Bloch mode which satisfies the conservation of momentum. If higher Bragg orders would be possible in the chosen range of parameters, we would have to analyze an individual band structure cut for each Bragg order.

In the chosen example, there are ten intersection points which come in pairs connected by a symmetry operation and describing counter–propagating waves. Thus, there are only five possible modes that must be further investigated and which I labeled with the letters 'a' through 'e' in Fig. 4.10. Another very important feature of the band structure is the anticrossing between band 6 and 7 at $k_\perp \approx 0.16$ and $\omega \approx 0.64$. For larger angles of incidence it approaches the operational frequency and complicates the situation around 51 degrees.

Next, we can plot the phase maps of the lattice–periodic part of the Bloch modes at the intersection points (cf. Fig. 4.11). The modes at 'a', 'b' and 'e' cannot couple to the zeroth Bragg order because they have a full phase change over one lattice constant. However, the mode 'c' couples to the in–plane component of the magnetic field and mode 'd' couples to any field component. Furthermore, the phase relations between all components with flat phase patterns are predominantly real which means that the Bloch modes couple very much like linearly polarized states rather than elliptical or even circular ones. This is in good agreement with the fact that the opal film behaves like a birefringent material.

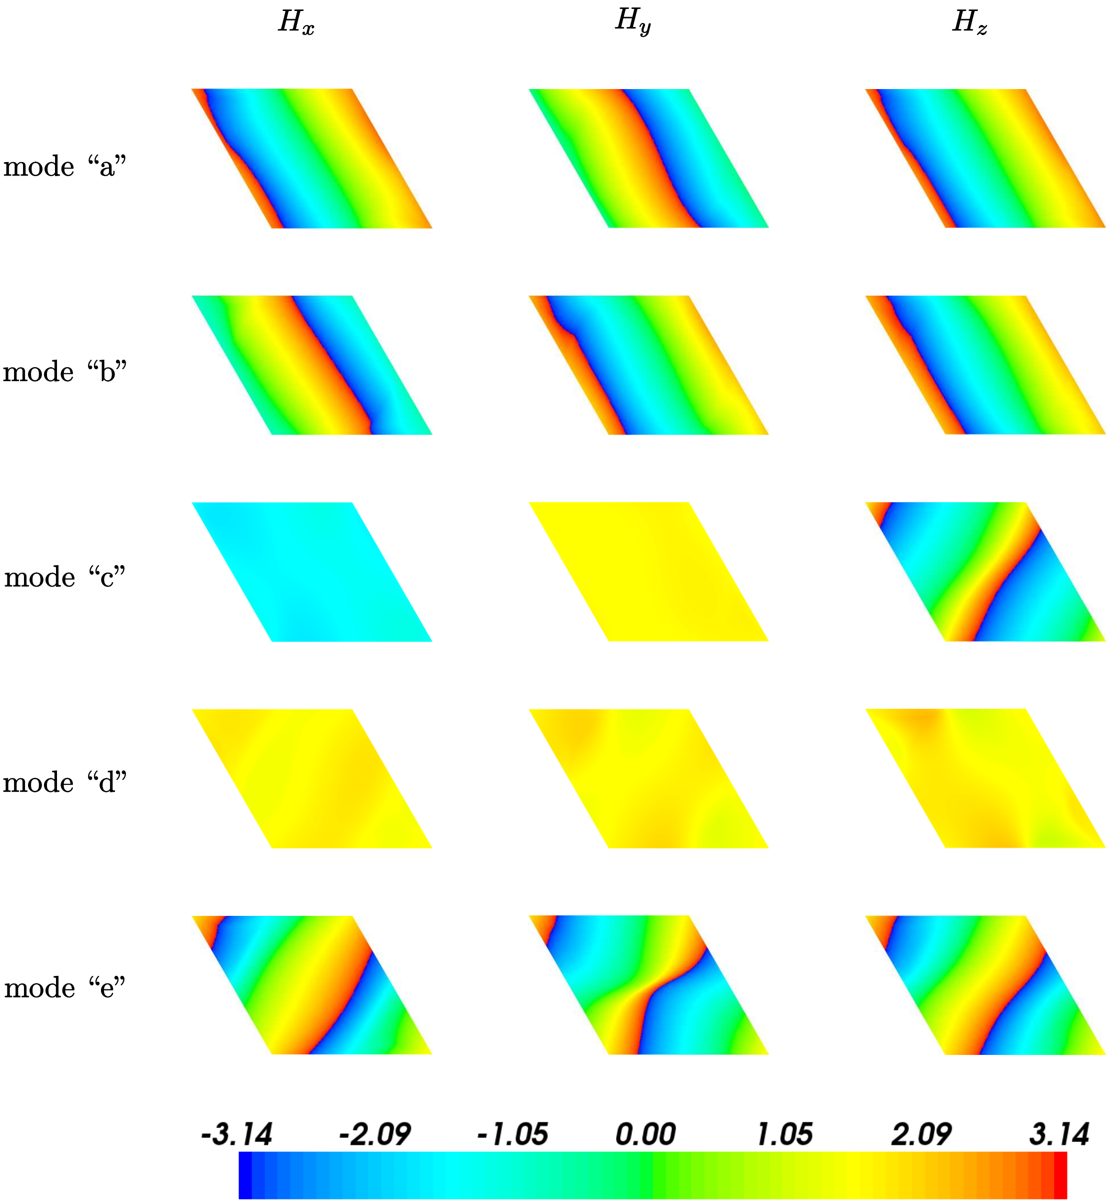

Figure 4.11.: *Phase plots of the five possibly coupling Bloch modes for an incident angle of 47 degrees (cf. Fig. 4.10). Each line corresponds to one Bloch mode. The columns depict the phases of the (from left to right) $x-$, $y-$ and z–component of the lattice–periodic parts. Clearly, only the modes ‘c’ and ‘d’ can couple to a zeroth order wave.*

Due to the anticrossing mentioned above, the incident wave can couple to the ‘e’–type mode for sufficiently large angles of incidence. This is clearly visible in Fig. 4.12.

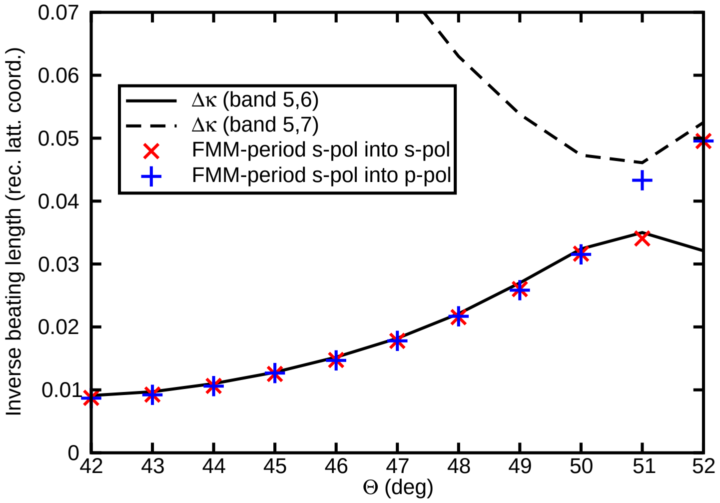

Figure 4.12.: *Comparison between the inverse beating periods (in inverse units of crystal planes) as obtained from the nonlinear fit of the FMM results (cf. Fig. 4.8) and from the appropriate band structure cuts (cf. Fig. 4.10). The incident wave has a wave length of* 900 *nm and impinges in the* Γ–L–W*–plane. For small angles, the agreement between the wave vector splitting and the beat length is very good. At 51 degrees, one FMM fit converges to the larger wave vector splitting whereas the other fit converges to the other splitting. In fact, both beat lengths form local minima of the nonlinear fit because the incident wave couples roughly equally to both Bloch modes that form this anticrossing.*

4.4.3. Comparison of FMM and Band Structure

Fig. 4.12 compares the polarization beat length derived from the FMM–calculations as well as predicted by the wave vector splitting found in the band structure. For angles below 50 degrees, both results nicely agree. At 51 degrees, the anticrossing is hit by the incident wave which consequently can couple to the three modes “c”, “d” and “e”. This results in a three–wave interference beat. Consequently, the fit ansatz as described above becomes a bit cumbersome and the accuracy of the fitted quantities is reduced. In fact, the nonlinear fit converges to either wave vector splitting (between band 5 and 6 and between band 5 and 7), depending on the starting point. A possible remedy would be to add a second sinusoidal term in Eqn. (4.15); unfortunately, this did not result in stable convergence of the fit minimizer.

At 52 degrees, the situation is relatively clear, again. However, fitting becomes more and more difficult because the beat amplitude decreases and approaches the noise level introduced by finite–size effects such as Fabry–Pérot resonances.

Finally, I compare the beat amplitudes of the cross– and bar–transmittance for various

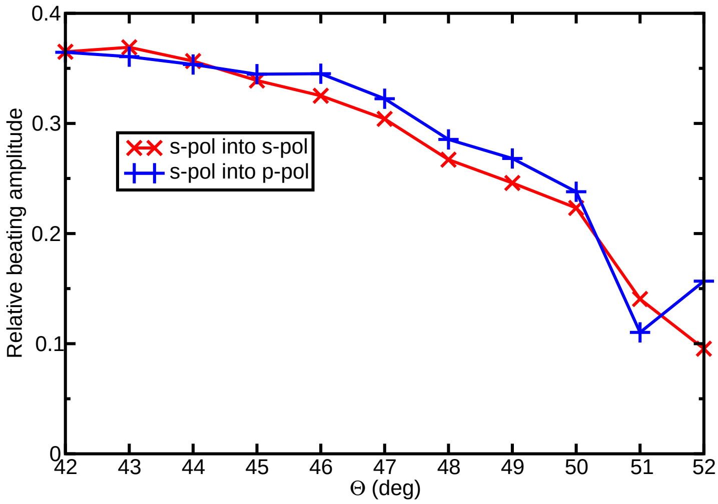

Figure 4.13.: *Comparison of the beat amplitudes of the cross–transmittance and bar–transmittance derived from the FMM fits (cf. Fig. 4.8). The fact that both curves are more or less on top of each other demonstrates that the PhC interface does not filter any polarization state. Again, the data become a bit shaky for larger angles because the nonlinear fit becomes difficult.*

angles (cf. Fig. 4.13) in order to find evidence for a polarization selectivity of the PhC interface. This would lead to a preferred transmittance into one of the polarizations. However, I could not find any strong effect which means that at least for these directions of incidence, the interface is insensitive to polarization. This is somewhat surprising because in combination with the plane of incidence, the interface does not feature any symmetries apart from one mirror plane.

4.4.4. Difference Between Experimental Data and Ab–Initio Numerics

The results obtained with the FMM are in good agreement with the band structure analysis performed above. Furthermore, they are qualitatively in good agreement with the experimental data. However, there are strong but very specific deviations that indicate a systematic deviation between the experimental samples and the simulated system. The quality of the samples looked very good in electron micrographs. Furthermore, unpolarized or bar–polarized long wave length transmission spectra were in excellent agreement with numerics. The two most prominent deviations, however, are a six–fold symmetry instead of a three–fold symmetry in the higher Bragg order transmittance spectra (not shown) and a strongly suppression of the cross–polarized transmittance into the zeroth Bragg order. In the latter case, it is particularly interesting that the bar–transmittance has the expected order of magnitude.

As the ab–initio calculations model many experimental details such as the substrate,

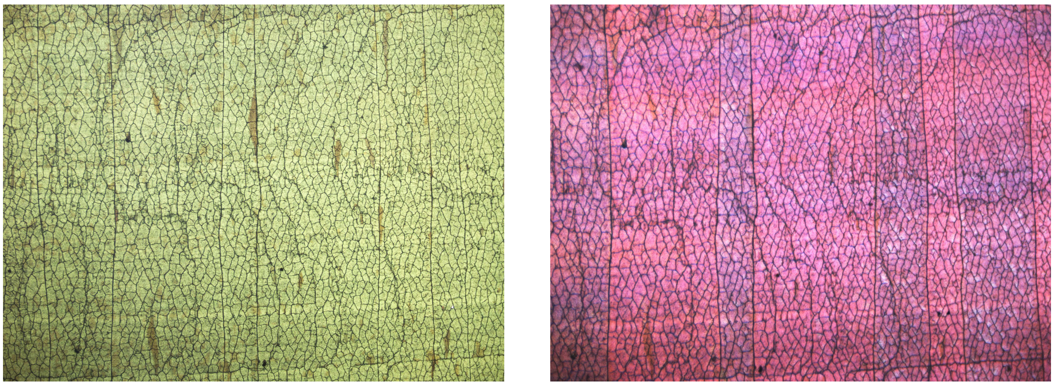

(a) unpolarized micrograph (b) polarized micrograph

Figure 4.14.: *Photographs of the opal film through a conventional microscope (left panel) and through a polarized microscope (right panel). The overall homogeneous color of the left picture indicates that the whole sample consists of fcc–opal. The right picture, however, reveals that the film consists of stripes with different polarization dynamics. Images are courtesy of Prof. Sergei G. Romanov.*

this difference must be due to some fabricational imperfection, e. g. , disorder or crystal defects. It is known [60] that opal films that are created by self–assembly are not necessarily monocrystaline. If, for instance, the parameters are not ideally chosen, the spheres might equally well arrange in other fashions than the desired face–centered cubic lattice. This results in a film that consists of differently stacked crystallites, which—due to the directional growth process—form stripes [60]. They are easily distinguished by means of their transmission or reflection spectra and are differently colored when the film examined through a microscope. Under such conditions, our samples were absolutely inconspicuous (cf. Fig. 4.14, left panel).

However, even for a perfect fcc–opal, there are two different modifications. The layers of spheres can be stacked either in the order "ABCABCA. . . " or in the reverse order "ACBACBA. . . ". If the polarization state of a light beam is modified upon transmission through a film, reciprocity requires that this change is undone if the transmitted ray is sent back through the film. Consequently, a polarization dynamics plot for an ACB–opal would be identical to Fig. 4.9 except that the circle on the Poincaré sphere would be passed through in the opposite direction.

Based on these results, I assumed that the experimental opal samples consist of stripes of ABC–opals and ACB–opals[4], whose bar–polarized transmitted light would interfere constructively in the far field, whereas the cross–polarized light would interfere destruc-

---

[4] I assumed that opal "nuclei" randomly form when the crystal growth starts and then are extended along the direction the meniscus moves over the substrate.

tively. Accordingly, I assumed that the detected cross–polarized transmittance signal is due to an imbalance between both types of opals. The easiest way to verify this is to look at the sample through a petrographic microscope. Furthermore, one would expect that the cross–polarized transmittance relative to the bar–polarized signal is decreased if the illuminated area of the film is increased.

When Prof. Romanov performed these checks, he indeed found differently colored stripes in the polarized micrographs, as depicted in Fig. 4.14. Furthermore, he found that increasing the spot size by a factor of 100 resulted in a reduction of the cross–polarized signal by 40%. This is less than what I expected and might indicate that the imbalance is relatively constant throughout the film. In other words, one type of opal seems to predominates in the sample, which is somewhat surprising.

5

Chapter 5.

Radiation Dynamics in a Woodpile Photonic Crystal

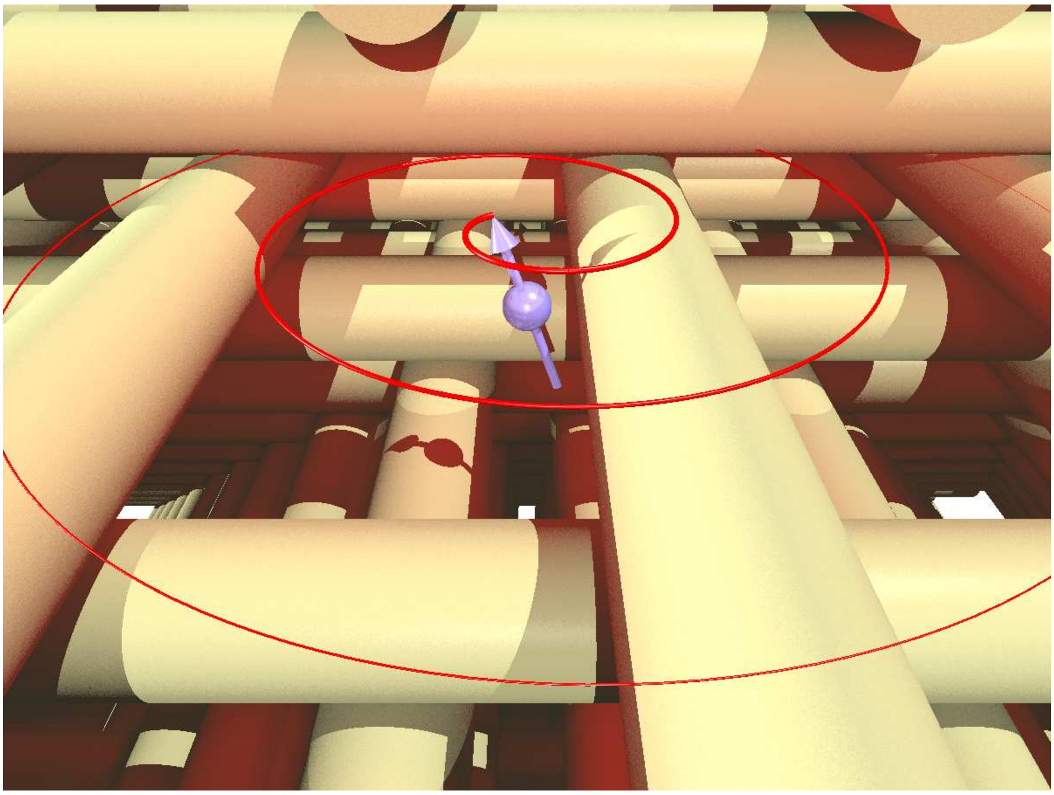

Illustration of a precessing dipole emitting radiation into a woodpile PhC.

The second practical project I would like to present in my thesis is my contribution to the investigation of the radiation dynamics in a photonic band gap material. I calculated the full memory kernel function of a radiating magnetic dipole in a (defect–free) woodpile structure. Other numerical results obtained with the Fourier Modal method (in particular the simulated transmission and reflection spectra of the finite sample) were calculated by Jens Küchenmeister who kindly allowed me to show them in this chapter. Finally, Prof. Ulrich Hoeppe and Prof. Hartmut Benner performed the experiment in Darmstadt and provided the experimental data.

This chapter starts with an introduction to this project (cf. Sec. 5.1). Then, I review the theory of spontaneous emission in a photonic crystal and adapt it to the case of a radiating magnetic dipole (cf. Sec. 5.2). In Sec. 5.3, I briefly present how the relevant quantities can be calculated numerically. After a discussion of some experimental details that make our approach particularly advantageous (cf. Sec. 5.4), I present the results (cf. Sec. 5.5), which lead to our manuscript Ref. [61]. The chapter ends with an outlook for further problems that could be investigated with the experimental method developed in Darmstadt.

5.1. Introduction

The dynamics of spontaneous emission is possibly the oldest topic in the context of photonic crystals because it motivated Eli Yablonovitch to propose photonic crystals in the first place [1]. Shortly thereafter, the emission dynamics in a band gap material was theoretically understood at least qualitatively [17, 62]. One of the main results was the prediction of strong non–Markovian effects such as Rabi–like oscillations and a giant Lamb shift close to a band gap edge. However, the theoretical models to represent a photonic crystal were somewhat crude and experimental work on this topic has been challenging.

I would like to emphasize that this chapter deals with the radiation dynamics of an emitter situated in a defect–free PhC[1]. I do not focus on emission into a photonic crystal cavity or the interaction between the emitter and otherwise (possibly due to disorder) localized modes. Plenty of work has already been done in this direction [26, 63].

Many experiments have been performed to investigate the radiation dynamics in a PhC most of which are performed with light. I would like to mention two popular approaches. In one class of experiments, quantum dots or dye molecules are introduced into a photonic crystal and excited by a laser pulse. Afterward, the decay characteristics of the spontaneous emission is recorded in time–domain and checked for non–exponential features [64, 65]. However, this approach does not allow to measure the anomalous Lamb shift.

In another class of experiments, narrow–band emitters such as rare–earth atoms are deposited in a PhC and the fluorescence spectrum is recorded with a spectrometer [66].

---

[1] Of course, the PhC in the experiment is of finite size which implies massive deviations from the infinite crystal. However, this cannot be avoided.

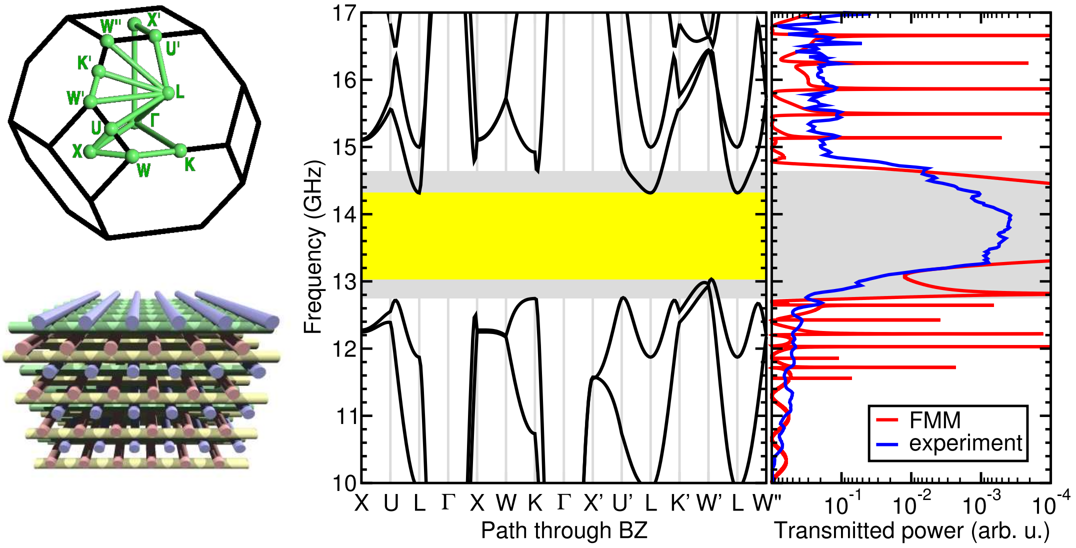

Figure 5.1.: *First Brillouin zone (leftmost), a band structure calculation (left part of the plot, calculated using the MPB [16] package) and transmittance data (right part of the plot) for our sample photonic crystal (woodpile of cylindrical rods as sketched in the leftmost column, $\epsilon = 9.8$; color was added for better contrast and is not physically motivated). The measured (blue line) and simulated (red curve; calculated by Jens Küchenmeister using the Fourier Modal Method [11, 12]) transmission spectra were obtained along the Γ–K direction and nicely reproduce the photonic stop band (gray shaded frequency range) in this direction. The complete photonic band gap is highlighted in yellow.*

This approach provides information about the Lamb shift because the transition frequencies of the atoms are well–defined. At the same time, however, this is a disadvantage because it restricts the spectral resolution to a few fixed frequencies. Furthermore, it is very difficult to detect the predicted time–domain effects. To summarize: Neither of the established methods at optical frequencies allows to measure the *full* radiation dynamics of a photonic crystal over a broad frequency range with good spectral resolution.

In this collaboration, we compare accurate measurements of the spontaneous emission line width and of the Lamb shift over a spectral range of more than an octave to ab–initio numerics. To our knowledge, this is the first measurement of the full linear radiation dynamics of a 3D photonic crystal in one experiment over a wide spectral range. Furthermore, we observed non–Markovian radiation dynamics.

To this end, Prof. Hoeppe performed an analogy experiment which was suggested by Prof. Busch. The basic idea was to model an excited atom (which often can be approximated as a two level system) by an appropriate spin system in a strong, constant magnetic field. This general approach to model quantum–optical processes in an easily

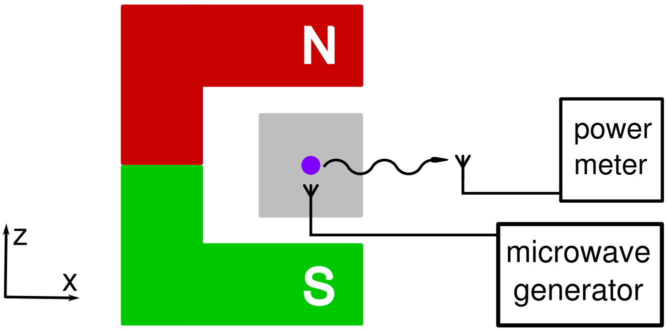

Figure 5.2.: *Basic sketch of the experimental setup. The YIG–sphere (violet) is situated inside the photonic crystal (gray) and between the poles of a magnet (red and green). The microwave excitation is generated by an HF–synthesizer and transmitted to the sphere via a small, nearby loop–antenna. The emitted radiation leaves the crystal, is received by a horn antenna and detected. Both antennae are oriented perpendicular to each other in the x–y–plane such that direct coupling is minimized. Unavoidable residual crosstalk is eliminated by an additional filter network (not shown, cf. Sec. 5.4).*

accessible frequency range had been suggested some fifteen years ago [67]. Nuclear magnetic resonance (NMR) was excluded due to the low frequency range. Electron spin resonance (ESR) occurs in a convenient spectral range but the signal to noise ratio turned out to be unsatisfactory. The final choice was a small sphere of yttrium–iron–garnet (YIG), which exhibits strong ferromagnetic resonance and is often used for tunable filters and oscillators in the microwave regime.

Using a YIG–sphere has many practical advantages (e. g. strong signal and convenient frequency range) but suffers from a fundamental flaw. In contrast to NMR and ESR samples, the response of a ferromagnetic resonator does not originate from independent spins but from collective spin excitations. These magnons are bosons. However, we restrict ourselves to the popular limit of a single excitation in the whole system (i. e. atom in the ground state and one photon in the crystal or no photons and excited atom). Within this limit, there is no difference between fermions and bosons and the YIG–sphere is an excellent model for a single spin in a magnetic field.

The model 3D photonic crystal is a woodpile structure with cylindrical rods (diameter $3.05 \pm 0.05\,\mathrm{mm}$) made of alumina ($\epsilon = 9.8$ according to the manufacturer). The lateral rod spacing was $10.6 \pm 0.03\mathrm{mm}$ from center to center. This structure was not at all optimized for the task of observing non–Markovian radiation dynamics but chosen according to Ref. [68]. According to band structure calculations (cf. Fig. 5.1), this structure exhibits a complete photonic band gap between 13.0 GHz and 14.3 GHz. To verify this, transmission spectra along several high–symmetry directions were measured and according simulations of the real (finite) structure were performed using the Fourier

Modal Method [11, 12], which is an ab–initio Maxwell solver. They are in good agreement with each other and the corresponding stop bands derived from the band structure (one such comparison is depicted in Fig. 5.1). A very rough sketch of the final setup is shown in Fig. 5.2.

5.2. Review of the Theory

The theory for the spontaneous emission of an atom in a non–trivial photonic environment has been known for over a decade [17, 69, 70] and typical non–Markovian effects such as an anomalous Lamb shift of an atom inside a band gap material have been predicted even earlier [62]. The atom is modeled as a two–level system that interacts with the light field via an electric dipole moment. In our experiment, however, we use collective excitations of electronic spins in an external magnetic field as a model for an atom and, consequently, we encounter magnetic dipole interaction. Although this should behave qualitatively similar to the case of electric dipole interaction, quantitative predictions can only be derived from a revised theory.

I refrain from presenting the calculation in detail as it does not provide too much additional knowledge as compared to the literature. Instead, I discuss the main differences due to the magnetic interaction and provide expressions that correspond to the main results of Ref. [70]. The derivations are analogous.

Throughout this discussion, I will use a notation which matches the rest of this thesis:

- ω (denoted Ω in Ref. [70]) is the operational frequency, i. e. the frequency at which radiation is observed.
- Ω (ω_{21} in Ref. [70]) is the naked resonance frequency of the atom, i. e. without coupling to any electromagnetic modes.
- $\tilde{\Omega}$ ($\tilde{\omega}_{21}$ in Ref. [70]) is the dressed resonance frequency of the atom, i. e. including the Lamb shift.
- $\alpha = \{n, \mathbf{k}\}$ (μ in Ref. [70]) is a composite index for labeling Bloch modes with wave vector $\mathbf{k}$ and band index n.

5.2.1. Normalization

First, we have to quantize the light field inside the crystal. To this end, we express the vector potential operator $\hat{\mathbf{A}}$ in terms of the classical vector potential $\mathbf{A}$ and the usual bosonic creation and annihilation operators $\hat{a}^\dagger$, $\hat{a}$. We choose to work in the Heisenberg picture and find

$$\hat{\mathbf{A}}(t) = \sum_\alpha \sqrt{\frac{\hbar}{2\epsilon c \omega_\alpha}} \left(\mathbf{A}_\alpha(t) \hat{a}_\alpha(t) + \mathbf{A}^*_\alpha(t) \hat{a}^\dagger_\alpha(t) \right) , \tag{5.1}$$

where the sum runs over all eigenmodes of the system. From this, we derive the magnetic field operator by applying the curl

$$\hat{\mathbf{H}}(t) = \sum_{\alpha} \sqrt{\frac{\hbar}{2\epsilon c \omega_\alpha}} \left[\big(\nabla \times \mathbf{A}_\alpha(t)\big)\, \hat{a}_\alpha(t) - \big(\nabla \times \mathbf{A}_\alpha(t)\big)^* \hat{a}^\dagger_\alpha(t) \right] . \tag{5.2}$$

As we would like to follow the lines of [69, 70], we normalize the vector potential accordingly[2]:

$$\int_{\mathrm{WSC}} \mathrm{d}^3 r \big(\nabla \times \mathbf{A}_\alpha(\mathbf{r})\big)^* \underbrace{\big(\nabla \times \mathbf{A}_{\alpha'}(\mathbf{r})\big)}_{\mathbf{H}^{(\mathrm{can})}_\alpha} = \frac{\omega_\alpha^2}{c^2} \delta_{\alpha\alpha'} . \tag{5.3}$$

Here, $\mathbf{H}^{(\mathrm{can})}$ is the magnetic field of the canonically normalized eigenmode. In vacuum, these modes would be plane waves and the magnetic field operator could be analytically expressed (cf. e. g. Ref. [71]). Here, this is not possible because the eigenmodes are not known analytically and must be determined via band structure calculations. These numerically obtained modes are usually normalized such that they fulfill:

$$\int_{\mathrm{WSC}} \mathrm{d}^3 r \Big(\mathbf{H}^{(\mathrm{num})}(\mathbf{r})\Big)^* \Big(\mathbf{H}^{(\mathrm{num})}(\mathbf{r})\Big) = \delta_{\alpha\alpha'} . \tag{5.4}$$

Consequently, the magnetic field operator can be constructed from the numerical Bloch modes $\mathbf{H}^{(\mathrm{num})}_\alpha$ by multiplying them with an additional factor ω:

$$\hat{\mathbf{H}} = \mathrm{i} \sum_{\alpha} \sqrt{\frac{\hbar \omega_\alpha}{2\epsilon c^3}} \left[\Big(-\mathrm{i}\mathbf{H}^{(\mathrm{num})}_\alpha\Big) \hat{a}_\alpha - \Big(-\mathrm{i}\mathbf{H}^{(\mathrm{num})}_\alpha\Big)^* \hat{a}^\dagger_\alpha \right] . \tag{5.5}$$

5.2.2. The Magnetic Dipole

The coupling between the orbiting electron in an atom and the electromagnetic field is relatively complicated and must be approximated to yield a decent interaction term. There are two popular approaches. Either the electron is approximated by a point dipole operator $\hat{\mathbf{d}}$ that couples to the electric field, leading to an interaction term of the form $-\hat{\mathbf{d}} \cdot \hat{\mathbf{E}}$, or it is modeled as a charged particle with momentum $\hat{\mathbf{p}}$ in a vector potential, leading in first approximation to an interaction term of the form $\hat{\mathbf{p}} \cdot \hat{\mathbf{A}}$. These approaches are inequivalent and yield results that are consistent on resonance but differ off resonance—more specifically, one factor ω in the former case is replaced by a factor Ω in the latter case.

In contrast, the localized spin resting[3] at a position $\mathbf{r}_0$ has a magnetic dipole moment $\mathbf{m}$, which couples to the magnetic field $\mathbf{H}$. Consequently, the proper interaction term in our case is a dipole coupling of the form $-\mathbf{m} \cdot \mathbf{H}$.

---

[2] Equation (A19) in Ref. [69]

[3] This is an approximation because the electrons in the YIG–sphere are not at rest. However, the additional interaction term due to spin–orbit coupling was expected to be small.

The magnetic dipole moment can be decomposed into two parts: First, a constant part along the magnetic field direction (canonically along the z–axis), which creates a static magnetization of the sample. The second contribution rotates in the perpendicular plane (x–y–plane) and describes the precession of the spin. This is the relevant quantity for modeling the radiation dynamics:

$$\hat{\mathbf{m}} = \boldsymbol{\mu}\hat{\sigma}_- + \boldsymbol{\mu}^*\hat{\sigma}_+, \tag{5.6}$$

$$\boldsymbol{\mu} = \mu\mathbf{e}_\mu. \tag{5.7}$$

If the magnetic moment belonged to a bosonic system (e. g. a magnon field), the Pauli operators would have to be replaced with bosonic creation and annihilation operators. However, this distinction is irrelevant for our discussion because we restrict ourselves to the single excitation case (i. e. an excited atom in a dark crystal or a ground state atom and exactly one photon in the crystal), anyway.

Consequently, the full Hamiltonian of the system reads

$$\hat{\mathcal{H}} = \sum_\alpha \hbar\omega_\alpha \hat{a}^\dagger_\alpha \hat{a}_\alpha + \boldsymbol{\mu}\mathbf{H}^{(\mathrm{ext})}\sigma_z + \hat{\mathcal{H}}_{\mathrm{int}}, \tag{5.8}$$

where $\mathbf{H}^{(\mathrm{ext})}$ is the external magnetic field which causes the precession. $\hat{\mathcal{H}}_{\mathrm{int}}$ is the interaction term

$$\hat{\mathcal{H}}_{\mathrm{int}} = \hat{\mathbf{m}} \cdot \hat{\mathbf{H}}(\mathbf{r}_0), \tag{5.9}$$

which reads within the rotating wave approximation (RWA)

$$\hat{\mathcal{H}}_{\mathrm{int}} = \mathrm{i}\hbar \sum_\alpha g_\alpha \hat{\sigma}_- \hat{a_\alpha}^\dagger - g^*_\alpha \hat{\sigma}_+ \hat{a_\alpha}, \tag{5.10}$$

where

$$g_\alpha = \mathrm{i}\sqrt{\frac{\omega_\alpha}{\hbar\epsilon c^3}}\boldsymbol{\mu} \cdot \left(\mathbf{H}^{(\mathrm{num})}_\alpha(\mathbf{r}_0)\right)^*. \tag{5.11}$$

This qualitatively differs from the coupling factor for electric dipole coupling and is the first main difference to Ref. [70]. The static external magnetic field is excluded in this last expression because it is perpendicular to the rotating contribution of the dipole.

5.2.3. Radiation Green's Function

The central quantity in the (linear) theory of spontaneous emission is the radiation Green's function. It describes the evolution of the probability that the atom is found in the excited state. At a given time, let the state of the two–level system be a superposition of its two eigenstates

$$|\,\Psi_{\mathrm{atom}}\,\rangle = a(t)|\downarrow\rangle + b(t)|\uparrow\rangle, \quad |a(t)|^2 + |b(t)|^2 = 1. \tag{5.12}$$

The projection of this state on the excited eigenstate evolves in time according to

$$b(t) = -\int_0^t \mathrm{d}t' \mathcal{G}_{\mathrm{rad}}(t-t')b(t'), \tag{5.13}$$

where $\mathcal{G}_{\text{rad}}(t-t')$ subsumes all possible events of emission of light, distributed scattering of the emitted light in the crystal and re–absorption by the atom [70].

Using the equations from Sec. 5.2.1 and Sec. 5.2.2, we can directly follow the lines of Ref. [70] to derive this function for the case of a radiating magnetic dipole. The radiation Green's function in a frame corotating with the bare atomic time evolution $\exp\big(\mathrm{i}\Omega t\big)$ reads:

$$\begin{aligned}\mathcal{G}_{\text{rad}}(\boldsymbol{\mu},\mathbf{r}_0,\tau) = \Theta(\tau)\,\beta \int_0^\infty \mathrm{d}\omega\;\omega\;\exp\big(-\mathrm{i}(\omega-\Omega)\tau\big)&\\ \times\sum_\alpha \delta(\omega-\omega_\alpha)\left|\mathbf{e}_\mu\cdot\mathbf{H}_\alpha^{(\text{num})}\right|^2.&\end{aligned} \tag{5.14}$$

This can be further simplified by inserting the expression for the projected local density of states in the magnetic field formulation

$$N_p^{(\text{num})}(\boldsymbol{\mu},\mathbf{r}_0,\omega) = \sum_\alpha \delta(\omega-\omega_\alpha)\left|\mathbf{e}_\mu\cdot\mathbf{H}_\alpha^{(\text{num})}(\mathbf{r}_0)\right|^2, \tag{5.15}$$

which we already defined in Sec. 2.3.1:

$$\mathcal{G}_{\text{rad}}(\boldsymbol{\mu},\mathbf{r}_0,\tau) = \Theta(\tau)\,\beta \int_0^\infty \mathrm{d}\omega\;\omega\;\exp\big(-\mathrm{i}(\omega-\Omega)\tau\big)\;N_p^{(\text{num})}(\boldsymbol{\mu},\mathbf{r}_0,\omega). \tag{5.16}$$

5.2.4. Angular Averaging

At this point, we encounter the second main difference between the case of a radiating atom and a radiating spin. Eqn. (5.14) contains the dipole orientation vector $\mathbf{e}_\mu$. As an atom is isotropic, there is no preferred direction of emission per se. However, each direction of emission corresponds to an individual decay behavior. Typically, this is solved by averaging the radiation Green's function over the full solid angle [70]. As the dipole orientation only occurs in the projected local densidy of states (LDOS), this is the only quantity to be averaged.

In contrast, a spin precessing in an external magnetic field (assumed to be applied in z–direction) has a well–defined dipole moment. It consists of a static contribution in z–direction and a dipole rotating in the x–y–plane:

$$\boldsymbol{\mu} = \mu(\hat{e}_x + \mathrm{i}\hat{e}_y). \tag{5.17}$$

Consequently, no averaging is necessary in our case. The LDOS simply reads:

$$N_p^{(\text{num})}(\boldsymbol{\mu},\mathbf{r}_0,\omega) = \sum_n \int_{\text{BZ}} \mathrm{d}^3k\;\delta(\omega-\omega_{n\mathbf{k}})\underbrace{\left|\mathbf{e}_\mu\cdot\mathbf{H}_{n\mathbf{k}}^{(\text{num})}(\mathbf{r}_0)\right|^2}_{P_{n\mathbf{k}}}. \tag{5.18}$$

To illustrate the correctness of this expression, we perform a consistency check. If the photonic system is invariant under time reversal (such that the rotational direction of the precession does not matter) and the starting orientation of the dipole is not specified, one would expect that the correct expression is equivalent to averaging over all orientations that the rotating dipole covers.

Time reversal symmetry implies

$$\omega_{n(-\mathbf{k})} = \omega_{n\mathbf{k}}, \tag{5.19}$$

$$\mathbf{H}_{n(-\mathbf{k})} = \mathbf{H}^*_{n\mathbf{k}}. \tag{5.20}$$

Eqn. (5.19) allows us to perform the k–space integration in Eqn. (5.18) twice

$$N_p(\boldsymbol{\mu}, \mathbf{r}_0, \omega) = \frac{1}{2}\sum_n \int_{\mathrm{BZ}} \mathrm{d}^3k \; \delta(\omega - \omega_{n\mathbf{k}}) \Big(P_{n\mathbf{k}} + P_{n(-\mathbf{k})} \Big) \Big). \tag{5.21}$$

Using Eqn. (5.20), we find

$$\begin{aligned} P_{n\mathbf{k}} + P_{n(-\mathbf{k})} &= \Big| H^{(x)}_{n\mathbf{k}} + \mathrm{i} H^{(y)}_{n\mathbf{k}} \Big|^2 + \Big| H^{(x)}_{n(-\mathbf{k})} + \mathrm{i} H^{(y)}_{n(-\mathbf{k})} \Big|^2 && (5.22) \\ &= \Big| H^{(x)} \Big|^2 + \Big| H^{(y)} \Big|^2 + \mathrm{i}\ \mathrm{Im}\Big\{ H^{(x)} H^{(y)*} \Big\} \\ &\quad + \Big| H^{(x)} \Big|^2 + \Big| H^{(y)} \Big|^2 - \mathrm{i}\ \mathrm{Im}\Big\{ H^{(x)} H^{(y)*} \Big\} && (5.23) \\ &= 2\Big| H^{(x)} \Big|^2 + 2\Big| H^{(y)} \Big|^2, && (5.24) \end{aligned}$$

where $H^{(x/y)}$ denotes the x– and y–components of the magnetic field, respectively. Indeed, this is the average over the coupling between the magnetic field pattern and all dipole orientations in the x–y–plane (modulo a prefactor). However, Eqn. (5.18) remains valid for magnetic systems and can be easily evaluated if the dipole rotates in a skew plane.

5.2.5. Emission Line and Lamb Shift

By inserting Eqn. (5.18) into Eqn. (5.16), we obtain the full radiation Green's function as a function of time. However, for comparison with experiments, we prefer a formulation in frequency–domain. To this end, we follow the lines of [70] to find the spectral representation of the Green's function

$$\begin{aligned} \mathcal{G}_{\mathrm{rad}}(\omega - \Omega) &= \gamma(\omega) + \mathrm{i}\lambda(\omega) && (5.25) \\ \gamma(\omega) &= \pi\beta\omega N_p^{(\mathrm{num})}(\omega) && (5.26) \\ \lambda(\omega) &= \frac{1}{\pi}\,\mathfrak{p}\!\!\int_0^\infty \mathrm{d}\omega' \frac{\gamma(\omega')}{\omega - \omega'}, && (5.27) \end{aligned}$$

where $\mathfrak{p}\!\!\int$ denotes the Cauchy principle value integral and $\beta = \frac{|\mu|^2}{8\pi^3 c^3 \hbar\epsilon}$. As compared to the result for electric dipole interaction, additional factors ω^2 and ω'^2 appear due to the different field normalization (Sec. 5.2.1).

From this, we can now explicitly find the spectral distribution of the atomic excitation

$$\tilde{b}(\omega-\Omega)=\frac{\mathrm{i}}{\omega-\Omega-\lambda(\omega)+\mathrm{i}\gamma(\omega)}. \tag{5.28}$$

According to the Wiener–Khintchine relation [72], the emission spectrum of the precessing dipole is essentially the real part of this excitation amplitude spectrum:

$$\begin{aligned} P_{\mathrm{emit}}(\omega) &= 2\mathrm{Re}\{b(\omega-\Omega)\} & (5.29)\\ &= P_0\,\frac{2\gamma(\omega)}{\big(\gamma(\omega)\big)^2+\big(\omega-\Omega-\lambda(\omega)\big)^2}. & (5.30)\end{aligned}$$

For a narrow line (i. e. γ and λ are approximately constant across the emission line) this is a Lorentzian with a line width of 2γ and whose maximum is shifted by λ. In the case of an atom, λ is the Lamb shift; in our case it is a classical analogon. It is connected to the line width via Eqn. (5.27) which closely resembles the Kramers–Kronig relations. The main difference is the lower integration boundary, which is $-\infty$ for Kramers–Kronig.

These terms were eliminated by the rotating wave approximation. In the full Hamiltonian in a Heisenberg picture, there are the slowly evolving coupling terms $\hat{a}\hat{\sigma}_+$ and $\hat{a}^\dagger\hat{\sigma}_-$ and there are fast evolving coupling terms $\hat{a}\hat{\sigma}_-$ and $\hat{a}^\dagger\hat{\sigma}_+$. The rotating wave approximation means that the latter terms are neglected [57]. In the Fourier transform these terms would cause another pole at $\frac{1}{\omega+\omega'}$, which would complete the integration domain. Unfortunately, the excitation number operator ceases to commute with the Hamiltonian if the fast evolving terms are taken into account and the calculation becomes much more tricky. In particular, it is no longer possible to stay in the single excitation limit.

5.3. Numerical Evaluation

In the previous section, I presented the theoretical framework for this project. However, there is no simple analytical solution of the wave equation for the photonic crystal at hand. Thus, I performed numerical band structure calculations. This section describes how the relevant quantities $\gamma(\omega)$ (the dispersive line width) and $\lambda(\omega)$ (the anomalous Lamb shift) are numerically derived from the band structure.

For the line width we need the local density of states at the YIG–sphere's position. To this end, I performed approximate on–shell calculations based on a linear **k**–space interpolation of the frequencies and eigenmodes as already presented in Sec. 2.3.3.

For calculating the Lamb shift, we have to evaluate a principal value integral

$$\lambda(\omega)=\frac{1}{\pi}\,\mathfrak{p}\!\int_0^\infty \mathrm{d}\omega'\,\frac{\gamma(\omega')}{\omega-\omega'} \tag{5.31}$$

for each target frequency ω. Here, $\gamma(\omega')$ was numerically calculated for a set of discrete frequencies $\{\omega_i\}$. As this calculation is quite expensive, the same frequency set has to be used for each integration and some sort of interpolation is unavoidable.

Numerical integration methods are not the first choice for this problem because the principal value integral in fact is a limit

$$\lambda(\omega) = \lim_{h\to+0} \frac{1}{\pi} \left[\int_0^{\omega-h} \mathrm{d}\omega' \frac{\gamma(\omega')}{\omega-\omega'} + \int_{\omega+h}^{\infty} \mathrm{d}\omega' \frac{\gamma(\omega')}{\omega-\omega'} \right], \tag{5.32}$$

which would be very prone to round–off errors because both terms diverge with opposite sign for $h \to 0$. As I am not very familiar with the peculiarities of numeric integration, I chose a different approach.

On each interval (ω_i, ω_{i+1}), we approximate $\gamma(\omega)$ by a linear interpolation of the numerically obtained values $\gamma(\omega_i)$ and $\gamma(\omega_{i+1})$. Accordingly, Eqn. (5.31) comes apart to analytically solvable sub–integrals over the respective frequency intervals. The integration scheme is presented in Appendix B. Conveniently, the resulting expression for such a sub–integral can be formulated such that it is valid irrespective of whether the pole lies inside the integration interval or not. This makes the implementation less prone to mistakes.

In this numeric calculation, both integration boundaries must be finite (fundamentally because the density of states is only known up to some cut–off frequency). Due to the first order pole of the integrand, the result is subject to a disturbance which decays logarithmically as the target frequency ω moves away from the cutoff. The resulting deviation from the correct result is very smooth and does not interfere with any local features of $\lambda(\omega)$.

The only fundamental restriction of this approach is that the integral must not be evaluated at one of the frequencies of support $\{\omega_i\}$. The relatively poor accuracy of the linear interpolation turned out to be well negligible as compared to numeric errors in the density of states and the logarithmic deviation introduced by the upper integration boundary.

5.4. $H^{(\mathrm{ext})}$–field Scans

A common feature of most experimental approaches in the optical regime is the fact that the resonance frequency of the quantum emitter is either fixed (e. g. transitions in an atom or a dye) or very difficult to control (e. g. quantum dots via temperature [73]). As a consequence, most experiments derive their information about the life time and Lamb shift from the optical spectrum [66] (hence from the frequency–domain) or from time–domain information [64]. There is a certain tuning range for quantum dots via the ambient temperature, which can be exploited to show strong coupling with a cavity [26]. However, I am not aware of any attempt to use this approach in the context of spontaneous emission in an ideal photonic crystal, possibly because the tuning range is very small.

However, a superficial inspection of the emitted power spectrum

$$P_{\mathrm{emit}}(\omega, \Omega) = P_0 \, \frac{2\gamma(\omega)}{\big(\gamma(\omega)\big)^2 + \big(\omega - \Omega - \lambda(\omega)\big)^2}, \tag{5.33}$$

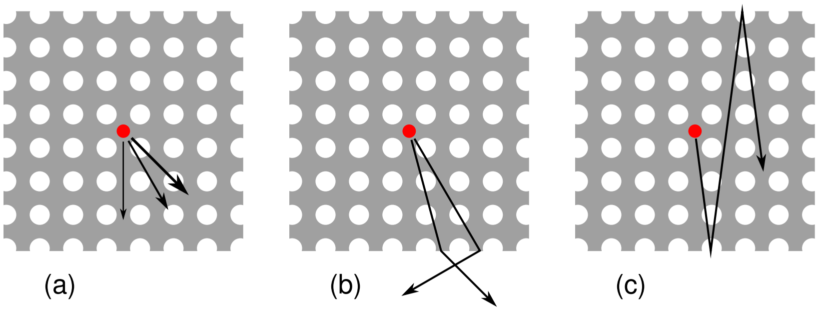

Figure 5.3.: *Some filter effects of a photonic crystal (gray structure) that affect the emission spectrum (i. e. radiated energy as a function of frequency) of the emitter (red dot) measured by a single, localized detector (not shown). Anisotropic emission of the emitter (a) occurs because Bloch modes with different wave vectors contribute differently to the projected LDOS. In time–resolved experiments, additional problems could occur because each mode travels with an individual group velocity. Unusual refraction at the interface (b) can diffract radiation towards the detector or away from it. Finally (c), some Bloch modes only couple very inefficiently to free space (see also Chap. 4). Such modes lead to very strong Fabry–Pérot resonances.*

(which has already been defined in Eqn. (5.30)) reveals that it would be very advantageous to keep the operational frequency ω (i. e. the one at which the emission is detected) fixed and vary the emitter's resonance Ω, instead: This equation describes a Lorentzian, whose parameters λ and γ depend on ω (but not on Ω). This ω–dependence distorts any frequency–domain spectrum and makes it impossible to measure λ and γ for a specific frequency.

Slightly more subtle errors are introduced into ω–domain data because the photonic crystal acts as a linear filter, as sketched in Fig. 5.3. Light propagation inside and refraction at the interfaces are highly anisotropic. As a consequence, the measured spectrum depends on the position and the orientation of the receiver[4]. This would not pose any problems if these effects were not highly dispersive, i. e. frequency–dependent. The resulting distortion of the line shape cannot be distinguished from a frequency–dependence of λ or γ and is expected to be strong in the vicinity of a stop band or a frequency with slow–light regime. Unfortunately, this is exactly the spectral range of interest. These problems can be systematically eliminated simply by measuring $P_{\text{emit}}(\Omega)$ instead of $P_{\text{emit}}(\omega)$. However, standing–wave–like effects due to a Bloch mode not being able to couple to propagating modes, modify the density of states inside the crystal and lead to periodic modifications in λ and γ, that cannot be easily compensated. Besides

---

[4] In optical experiments, this is less severe as compared to the microwave regime because the emission can be collected with a lens with high numerical aperture. This averages over a large number of directions.

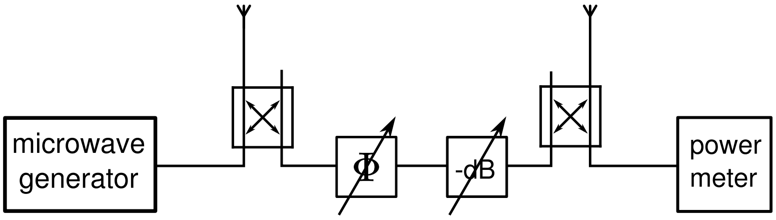

Figure 5.4.: *Schematic of the compensation network developed by Prof. Hoeppe to eliminate the residual direct coupling of the exciting loop and the receiving horn antenna: A portion of the HF signal is sidelined from the feeding line via a directional coupler, modified in phase and attenuated. For each frequency independently, the parameters are chosen such that they exactly compensate the direct irradiation when fed into the receiving line via another directional coupler. Each setting compensates the parasitic coupling for exactly one frequency. Thus, this type of compensation requires to perform magnetic field scans with ω kept fixed.*

this problem, it is possible to measure accurate values $\lambda(\omega)$ and $\gamma(\omega)$ for a single ω by recording the resonance as a function of Ω. The result is a perfect Lorentzian (cf. Fig. 5.5).

In principle, this is possible for many emitters such as quantum dots (controllable via temperature or Stark– or Zeeman shifts). The great advantage of a precessing magnetic dipole is that Ω can be accurately controlled over a wide spectral range via the static external magnetic field and—indeed—allows for measuring the resonance as a function of Ω.

Prof. Hoeppe decided to follow this path not because of the fundamental advantages stated above but for a very practical reason: The (symmetric) exciting loop antenna was fed by a (electrically asymmetric) coax line. The resulting sheath currents emitted radiation so strongly that the signal of the YIG–sphere was buried in this emission. To eliminate this parasitic coupling, he added a compensation network (cf. Fig. 5.4), which requires the operational frequency to be kept fixed.

5.5. Line Width and Dispersive Resonance Shift

In this section, I first present the results of my computations as well as the experimental findings. The data are shown for two distinct positions of the YIG–sphere inside the PhC as indicated by Fig. 5.6. Then, I estimate the numerical errors due to the finite discretization of the isofrequency integration and the finite frequency interval over which the principal value integration was performed.

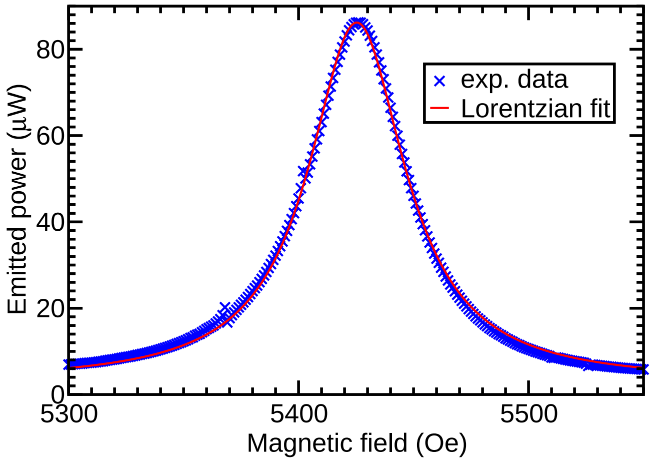

Figure 5.5.: *The experimentally measured resonance line as a function of the magnetic field with the operational frequency kept fixed at* 14.9 *GHz. Although this frequency is just above the band gap (where a strongly distorted resonance line would be expected) the magnetic field–scan is a near–perfect Lorentzian. This plot was adapted from Ref. [61].*

5.5.1. Comparison of Theory and Experiment

Prof. Hoeppe performed magnetic field scans for closely spaced frequencies for two distinct positions of the YIG–sphere. These spots are indicated in Fig. 5.6. From these data, he extracted the line width and resonance shift information, which is the difference between the actual resonance position and the expected magnetic field on resonance according to Ref. [74].

I calculated the same quantities for a point–dipole at the respective positions. Both, theoretical and experimental, data sets are plotted in Fig. 5.7. They completely determine the radiation Green's function throughout the covered frequency window. Hence, the full linear radiation dynamics has been characterized.

The agreement between theory and experiment is very good. As expected[5], the line width in the effective material limit (small frequencies) is approximately cubic and greater than the vacuum curve by about the mean dielectric constant of the crystal. Inside the band gap, any sort of emission is strongly suppressed. Finally, we can see a strong dependence of the line width on the sphere's position. In terms of the resonance shift, we can see that it is rather small in the effective limit and grows dramatically at the lower band gap edge. Inside the band gap, it experiences a steep decrease towards a minimum at the upper band gap edge.

All features described above are clearly visible both in the theoretical and the experimental data. However, there are several deviations that we believe to understand:

---

[5] The number of available wave vectors in a homogeneous material grows quadratically with ω and another factor of ω enters because of the normalization condition Eqn. (5.3).

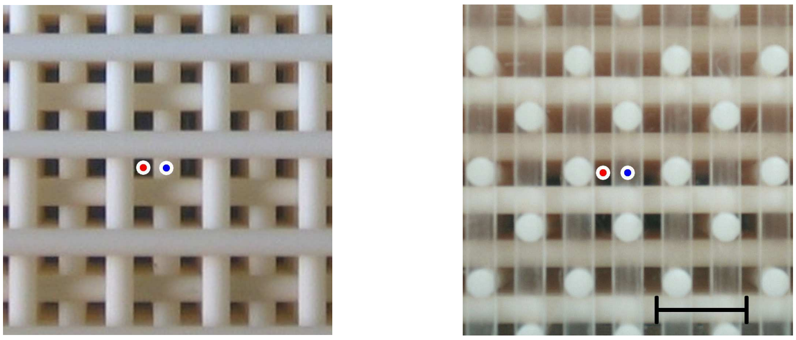

Figure 5.6.: *Top view (left panel) and side view (right panel) of the woodpile photonic crystal employed in this project. The bar corresponds to a length of* 10 *mm. Furthermore, this figure indicates the positions of the YIG–sphere for which results are presented in this chapter. Position "a" is next to a dielectric rod and indicated by the red dots. Position "b" is as far away from any dielectric as possible and indicated in blue. This figure was taken from our manuscript Ref. [61]. Copyright has been transferred to the the American Physical Society in case of publication.*

1. Strong oscillations of the line width just below the band gap are due to standing waves (Fabry–Pérot resonance) inside the photonic crystal sample but not covered by the band structure, which describes an infinite PhC. These oscillations are a fundamental problem of every finite sample and increase in number as the sample is increased in size. Thus, our relatively small sample has the advantage that the standing waves can be experimentally resolved.
2. Weaker oscillations above the band gap can be attributed to the same effect.
3. Due to its Kramers–Kronig–like connection to the line width, the resonance shift features a falling slope wherever the line width has a dip.
4. The theoretical resonance shift at higher frequencies is artificially reduced due to the finite integration domain (cf. Sec. 5.3).

The most obvious difference, however, is not fully understood: At position "b" and for frequencies above the band gap, the experimental line width is considerably smaller than expected. A possible explanation would be poor coupling between the exciting loop antenna and the sphere because the relevant Bloch mode has a zero at the sphere's position. Another reason could be an influence of the feeding coax line or the sphere's mount. Finally, a "caught" Bloch mode (i. e. one that cannot couple to plane waves outside the PhC) could contribute significantly to the projected LDOS of the infinite crystal. In the experiment, such a mode would not contribute to the line width as it cannot carry energy to infinity. None of these explanations is really satisfactory and—more important—there is no experimental evidence for any of them. Besides this, the

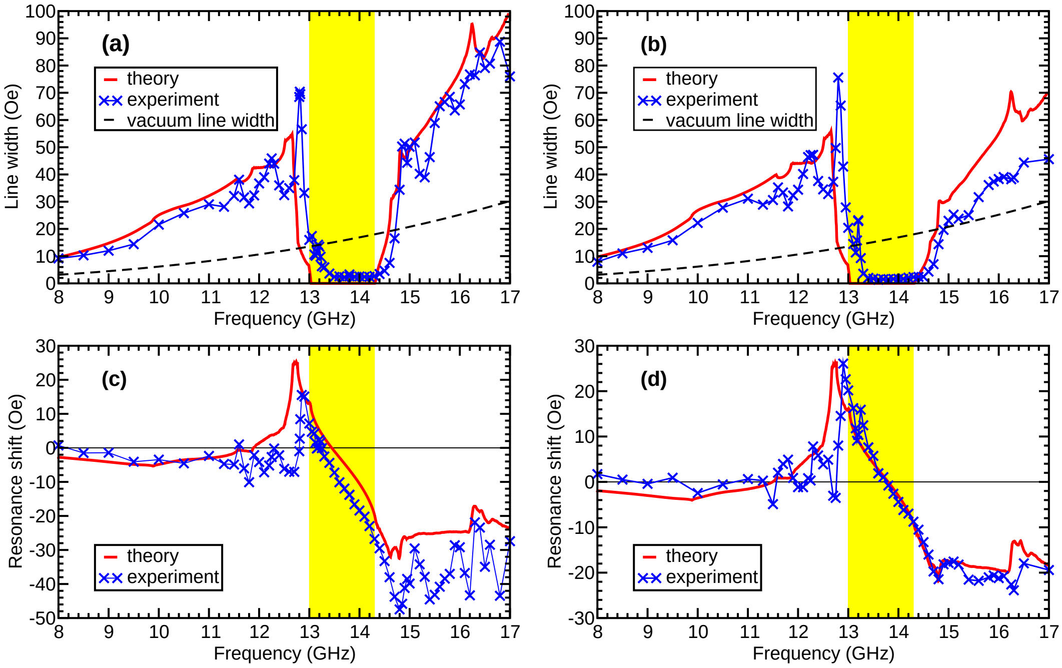

Figure 5.7.: *Dispersive line width and resonance shift plots for the two sphere positions illustrated in Fig. 5.6: position "a" in the left column and position "b" on the right hand side. The top row shows the theoretical line width, the measured one and the line width expected for emission in vacuum. The bottom row shows the theoretical and the measured resonance shift. All ordinates are in Oersted and can be translated into frequencies by multiplying with* 2.81 *MHz/Oe. The band gap (cf. Fig. 5.1) is highlighted in yellow. These plots were adapted from Ref. [61].*

agreement between theory and experiment is acceptable.

5.5.2. Non–Markovian Radiation Dynamics

Markovian dynamics is memory–free. Consequently, a non–Markovian process is characterized by a time evolution including a non–trivial memory kernel Green's function:

$$\partial_t x(\omega) = \int_{-\infty}^{0} \mathrm{d}t' \mathcal{G}(t-t')x(t'). \tag{5.34}$$

As a Dirac distribution in time–domain translates to a flat frequency–domain representation $\mathcal{G}_{\mathrm{rad}}(\omega)$, *any* frequency–dependence of the memory kernel implies non–Markovian dynamics in a strict sense. As the vacuum density of states grows quadratically, even vacuum (strictly speaking) is non–Markovian. However, the relevant question is whether $\mathcal{G}_{\mathrm{rad}}(\omega)$ is approximately constant over the spectral range covered by an emitter (i. e. its

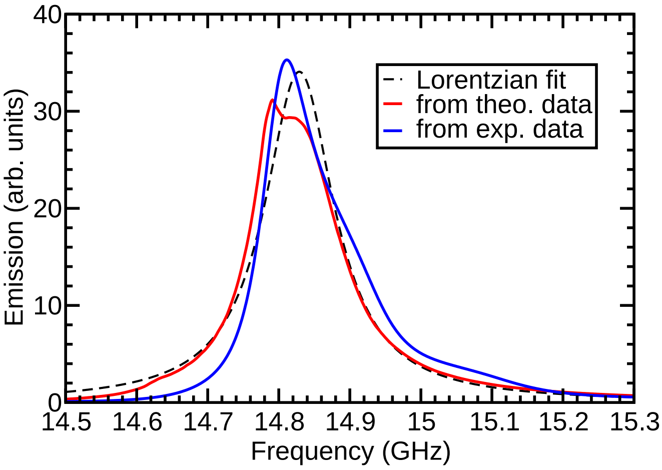

Figure 5.8.: *Reconstructed frequency–domain resonance lines with the magnetic field kept fixed. These lines were constructed from the spectral representation of the calculated and measured radiation Green's function. By this approach, I eliminated all linear filter action of the photonic crystal—the plots depict the total power emitted by the YIG–sphere. It is nicely visible how these frequency–domain lines are distorted due to the dispersive nature of* γ *and* λ *although the corresponding* $H^{(ext)}$*–field scan (cf. Fig. 5.5) is perfectly Lorentzian. For comparison, I added a Lorentzian fit to the experimental data (black dashed line). This plot was adapted from Ref. [61].*

line width). In such a case, the non–Markovian nature of the process is not visible due to the fast decay and the Wigner–Weisskopf approximation [75] is justified. This holds true for vacuum in all but rather exotic cases.

However, abrupt spectral features in the density of states such as the isolated resonance of a cavity or a photonic band gap occur in a frequency interval comparable to the line width. In such a case, the (ideally Lorentzian) emission line as a function of frequency is distorted.

To check whether such a non–Markovian regime exists in our problem, I evaluated Eqn. (5.30) using the theoretical and measured values for $\gamma(\omega)$ and $\lambda(\omega)$. The result is shown in Fig. 5.7. The resulting virtual emission lines at 14.9 GHz (i. e. just above the band gap) are depicted in Fig. 5.8. Although they differ rather strongly due to the Fabry–Pérot oscillations in the experiment, they both are skewed towards the band gap. This proves that—indeed—the radiation Green's function varies so rapidly with respect to ω that we may speak of non–Markovian radiation dynamics.

5.5.3. Numerical Accuracy

The theoretical line width plots shown in Fig. 5.7 were constructed by interpolating the band structure on a total of 2.4 million tetrahedra tessellating the first Brillouin zone. As I used a linear interpolation scheme, we may expect first order convergence with respect to the resolution in reciprocal space. Based on this and the results of coarser calculations (not shown), I estimate the accuracy of the numerical line width plot to be approximately 3.5%.

The theoretical resonance shift plots were derived from similar curves based on a coarser tessellation (relative accuracy 5%). However, they are subject to further errors. Although, the numerator of the integrand was piecewise linearly interpolated, this did not significantly contribute because the sampling was chosen very fine. Much more severe is the finite upper integration boundary, which causes a logarithmically decaying deviation. In order to reduce this effect, the integration was performed on a relatively large frequency window (0–30 GHz). As the number of bands (and, hence, the memory consumption) grows approximately cubically with ω, this could be done only using said relatively coarse tessellation. In total, the relative deviation[6] due to the integration cut–off is as high as 50%, yet very smooth without distorting local features.

5.6. Outlook: Possible Applications in Other Systems

At this point I would like to provide some speculations which other systems could be worth investigating with the method presented in this chapter. Besides the interaction of individual emitters and other nano–photonic systems such as metallic metamaterials, it could be used to experimentally study the statistical properties of random dielectric media. Under certain circumstances, such systems are predicted to exhibit Anderson localization, i. e. the existence of extended, yet exponentially localized eigenmodes. This results in an exponential dependence of the transmittance with respect to the thickness of the sample [76]. Unfortunately, absorption has a similar effect on the transmittance and it could be advantageous to measure a quantity which is differently influenced by absorption and localization.

As we have seen in this chapter, a single YIG–sphere can be used to accurately scan the radiation Green's function. This quantity is closely related to the photonic system's single particle Green's function $\mathcal{G}_{\mathrm{sp}}$ whose imaginary part is the local density of states:

$$N(\omega, \mathbf{r}) = \mathrm{Im}\{\mathcal{G}_{\mathrm{sp}}(\mathbf{r}, \mathbf{r}; \omega)\}. \tag{5.35}$$

Comparison with the predictions of analytical theory would be possible by performing such measurements for different realizations of the random system. The statistical

---

[6] This quantity is very prone to large relative errors in the experiment, too. This is because it is measured as the difference between two approximately equal numbers: The bare resonance frequency and the dressed one.

moments derived from these data could be compared to the disorder–averaged quantities that result from effective medium theory. I expect the statistical properties of a disordered system to be well–distinguishable from the effects of dissipation.

6

Chapter 6.

Conclusion

Now that I presented the findings of my studies, I would like to conclude this thesis with a brief summary of the preceding chapters and an outlook of interesting open questions.

6.1. Summary

In this thesis, I investigated several photonic crystal structures by means of the photonic band structure and related numerical tools.

In Chap. 2, I reviewed the popular approaches to solve both conventional and on–shell band structure problems. Besides that, I discussed the question of when a conventional band structure calculation on a set of high–symmetry lines can be safely assumed to show the correct band gap. The main result is that in highly symmetric structures, certain high symmetry points are guaranteed to exhibit stationary points. These points are natural candidates for the band extrema that define the band gap. Having said that, an investigation of tight–binding models suggests that long–range interaction between the constituent scatterers of the crystal as well as strong inter–band coupling increases the number of stationary points. Furthermore, I provide some rough rules of thumb that were helpful for my band structure calculations.

Chap. 3 pertains most of my work on the topic of photonic Wannier functions. My main contributions to the advancement of the Wannier method in photonics are a distinct improvement of the algorithm to create the Wannier basis and the further development of Wannier–based simulation tools. The former resulted in a reduction of the runtime for creating WFs by some orders of magnitude and the incorporation of additional constraints. In the latter context, I was able to show that absorbing boundaries in PhC wave guides imposed via a one–way wave equations are not efficient. Finally, I found discontinuities in the physical eigenmodes of every three–dimensional photonic crystal. This necessarily results in the poorly localized Wannier functions for those systems.

In Chap. 4, I performed a theoretical analysis of a particular polarization conversion effect that had been found experimentally for light transmitted through thin opal films under oblique incidence. To understand the underlying physical process, I performed a close analysis of ab–initio calculations of the transmission process and of the opal's band structure. This revealed that the polarization conversion is due to an interference beat between Bloch modes. The beat strongly resembles linear birefringence known from conventional crystal optics. However, opals exhibit the full cubic symmetry of the fcc–lattice, hence, optical anisotropy would conventionally not be expected. This clearly demonstrates that a description of PhCs using effective material parameters breaks down as soon as the long wave length regime is left. Finally, we found a quantitative discrepancy between the theoretically expected and experimentally measured cross–polarized transmission. I interpreted this result as a sign for the coexistence of differently stacked domains of fcc–opal in the sample. A subsequent investigation of the opal sample confirmed this hypothesis.

The dynamics of spontaneous emission has been of particular interest in the photonic crystal community since the first publications by Eli Yablonovich and Sajeev John [1, 2]. In Chap. 5, I present results on the emission dynamics of a single point–like emitter inside a band gap material. I compare my theoretical results to experimental data that were obtained by means of a ferromagnetic resonator inside a macroscopic woodpile structure, driven at microwave frequencies. This novel experiment allows to accurately measure the central quantity for radiation dynamics: the radiation Green's function. The theoretical and experimental results for two distinct emitter positions inside the crystal are in excellent agreement over a spectral range of more than one octave. They exhibit a strong suppression of emission inside the band gap and reveal the existence of a strong anomalous Lamb shift and of non–Markovian dynamics in this system.

6.2. Outlook

This thesis provides only a small step towards a complete understanding of photonic crystals. Some questions that may be interesting and lead to further insight into the properties of structured dielectric media were already formulated in the chapters where they naturally arise. Here, I would like to gather these topics that may be worth further investigation.

Stationary Band Structure Points

My discussion about the occurrence of stationary points in photonic band structures is rather sketchy and far from being exact. However, this topic is very interesting and could be worth further investigation. In particular, I refrained from closely analyzing the effects of interband coupling and only provided a very rough upper bound for the number of stationary points. Closely related to this is the problem of long–range interaction of non–isolated bands.

On–Shell Band Structure

The problem of calculating on–shell band structures is not efficiently solved, yet. There are two main problems: First, the eigenequation with respect to the wave vector is nonlinear and not trivially linearizable. Second, unphysical solutions cannot be excluded a priori. Interpolation techniques based on a conventional band structure calculation on a discretization of the first Brillouin zone seem to be among the most promising approaches. The linear interpolation scheme employed in this thesis is easy to implement but exhibits only first order convergence. More sophisticated interpolation schemes such as higher order polynomials have the inherent disadvantage that the occurring equations become much more complicated. In my opinion, a combination of linear interpolation and the Wannier basis (which effectively is a multi–dimensional Fourier–interpolation) would be worth investigation: The Wannier set[1] allows to formulate the on–shell problem as a highly nonlinear eigenproblem with an extremely small matrix. Starting with an approximate solution from the linear interpolation, it could be solved via Newton's method within only few iterations.

Symmetric Wannier Functions

Symmetries are a key feature of most crystals and of great practical value to reduce the complexity of physical problems. This holds true in the context of photonic Wannier functions as well. For instance, the degree of localization that can be obtained for a set of Wannier functions depends strongly on the hybridization; exponential localization can be obtained only if the Bloch modes' symmetries are properly taken into account [22].

Such exponentially localized Wannier functions would be of great practical value in numerical calculations. They would result in exponential convergence with respect to the long–range interactions taken into account. However, the spread minimizer itself is blind to symmetry. It is likely to converge towards functions with a slowly decaying tail if their spread is only slightly smaller than that of the exponentially localized ones. Thus, a symmetrized version of the spread minimization algorithm would be of great practical relevance.

Three possible approaches immediately come to mind:

1. Use Eqn. (3.51) to express the spread functional on only the irreducible wedge of the first Brillouin zone. This would reduce the number of degrees of freedom by a factor that is slightly smaller[2] than the number of elements in the PhCs point group and thus increase the performance and stability of the Wannier generation. Unfortunately, this approach leads to problems for correctly expressing the overlap

---

[1] The unitary matrices alone are sufficient for basic operation; the WFs themselves are required if the Bloch modes are sought–after as well.

[2] The $\mathbf{k}$–mesh is discrete and points that would belong only "partially" to the irreducible wedge must be added completely. This holds for all points that are mapped onto themselves by at least one symmetry operation, most notably Γ.

matrices $\mathcal{M}^{(\mathbf{k},\mathbf{b})}$ (cf. Eqn. (3.12)) on or next to border points of the irreducible wedge.

2. Use Eqn. (3.51) to symmetrize the gradient before feeding it to the standard WF generation. This would not reduce the number of degrees of freedom but still would yield exactly symmetric functions. A problem (that is also present for the previous approach) is that the representation matrices for the Bloch modes at all **k**–points have to be known. In principle, they can be numerically calculated, albeit at relatively high computational cost.
3. Use Eqn. (3.51) to derive an auxiliary penalty term that makes all asymmetric configurations so unfavorable that they are avoided by the minimizer. This approach is sketched in Sec. 3.3.7.

In my opinion, the last approach poses the least problems and should be tested first. In Sec. 3.3.7, I derived some possibly helpful equations.

Absorbing Boundaries in Discrete Wave Guides

Absorbing boundaries are a crucial prerequisite for the simulation of open systems. The wave guides discussed in Sec. 3.5.4 are related to other discrete and periodic problems in physics such as single–photon transport in coupled resonators. Thus, a method to efficiently impose open boundaries to PhC wave guides might be valuable for certain problems in solid–state physics and quantum optics as well.

In Sec. 3.5.4 I demonstrated that the approach of imposing absorbing boundaries via a one–way wave equation in general is not feasible in periodic systems. As already mentioned, the somewhat related approach of suppressing the reflected waves via absorption in a non–reciprocal wave guide seems to be more promising.

Further Interference Beats in Opal Films

The polarization interference beat presented in Chap. 4 exhibits many similarities to conventional birefringence. However, interference beats between Bloch modes result not only in effects that have analoga in crystal optics. For instance, they can cause an oscillatory dependence of diffraction into several Bragg orders[3]. Furthermore, the presented effect already features a three–wave interference beat, that is unknown in homogeneous media. Many of these interference effects are very sensitive to the propagation direction and frequency of the incident light, which makes them interesting for sensing applications such as the measurement of angles.

Emission in Structured Media

First, I would like to mention the possibility to use the experimental technique described in Chap. 5 to investigate related dielectric or metallic structures. This would provide a

---

[3] This effect exists, has been experimentally verified and is currently analyzed in our group.

unique way to identify strong localization in disordered media via the statistical properties of the local density of states as already mentioned in Sec. 5.6.

Second, the results shown in this chapter are convincing but some questions remained unanswered. The most important is the unexpectedly narrow experimental line width above the band gap. Furthermore, the photonic crystal structure used for this experiment was not chosen based on any estimation of how strong the non–Markovian dynamics would be. This opens the interesting question of how to optimize structures for less common figures of merit such as a steep drop of the density of states close to a band gap.

A

Appendix A.

Derivation of Secondary WF–Objectives

In this appendix, I provide intermediate steps for deriving the $\mathcal{U}^{(\mathbf{k}}$–gradients of the auxiliary penalty terms introduced in Sec. 3.3.6 and Sec. 3.3.7.

A.1. Imaginary Part

The imaginary part of a Wannier function is a real-valued function which reads (following the usual conventions):

$$\mathrm{Im}|\,\mathbf{W}_{n\mathbf{0}}\,\rangle = \frac{1}{2\mathrm{i}}\left(|\,\mathbf{W}_{n\mathbf{0}}\,\rangle - |\,\mathbf{W}_{n\mathbf{0}}{}^{*}\,\rangle\right). \tag{A.1}$$

Here, $|\,\mathbf{W}_{n\mathbf{0}}{}^{*}\,\rangle$ denotes the complex conjugated of the Wannier function $|\,\mathbf{W}_{n\mathbf{0}}\,\rangle$.
We aim to express the norm of the imaginary part of a single WF:

$$\begin{aligned}
\Omega^{(n)}_{\mathrm{imag}} &= \Big|\,\mathrm{Im}|\,\mathbf{W}_{n\mathbf{0}}\,\rangle\,\Big|^{2} && \text{(A.2)}\\
&= \frac{1}{4}\Big[\left(\langle\,\mathbf{W}_{n\mathbf{0}}\,| - \langle\,\mathbf{W}_{n\mathbf{0}}{}^{*}\,|\right)\left(|\,\mathbf{W}_{n\mathbf{0}}\,\rangle - |\,\mathbf{W}_{n\mathbf{0}}{}^{*}\,\rangle\right)\Big] && \text{(A.3)}\\
&= \frac{1}{4}\left[2\langle\,\mathbf{W}_{n\mathbf{0}}\,|\,\mathbf{W}_{n\mathbf{0}}\,\rangle - \langle\,\mathbf{W}_{n\mathbf{0}}{}^{*}\,|\,\mathbf{W}_{n\mathbf{0}}\,\rangle - \langle\,\mathbf{W}_{n\mathbf{0}}\,|\,\mathbf{W}_{n\mathbf{0}}{}^{*}\,\rangle\right] && \text{(A.4)}\\
&= \frac{1}{2}\Big[1 - \mathrm{Re}\left\{\langle\,\mathbf{W}_{n\mathbf{0}}{}^{*}\,|\,\mathbf{W}_{n\mathbf{0}}\,\rangle\right\}\Big], && \text{(A.5)}
\end{aligned}$$

in terms of the unitary mixing matrices. To this end, we must reformulate the overlap product in Eqn. (A.5). First, we project onto the Bloch modes and the complex conjugated Bloch modes, respectively:

$$\begin{aligned}
\langle\, \mathbf{W}_{n\mathbf{0}}{}^{*} \,|\, \mathbf{W}_{n\mathbf{0}} \,\rangle =& \frac{1}{V_{\mathrm{BZ}}} \sum_{m} \int_{\mathrm{BZ}} \mathrm{d}^3k \, \langle\, \mathbf{W}_{n\mathbf{0}}{}^{*} \,|\, \mathbf{H}_{m\mathbf{k}} \,\rangle \langle\, \mathbf{H}_{m\mathbf{k}} \,|\, \mathbf{W}_{n\mathbf{0}} \,\rangle \\
=& \frac{1}{V_{\mathrm{BZ}}^2} \sum_{m,l} \int_{\mathrm{BZ}} \mathrm{d}^3k \int_{\mathrm{BZ}} \mathrm{d}^3k' \, \langle\, \mathbf{H}_{l\mathbf{k}'} \,|\, \mathbf{W}_{n\mathbf{0}} \,\rangle \\
& \times \langle\, \mathbf{H}_{l\mathbf{k}'}{}^{*} \,|\, \mathbf{H}_{m\mathbf{k}} \,\rangle \langle\, \mathbf{H}_{m\mathbf{k}} \,|\, \mathbf{W}_{n\mathbf{0}} \,\rangle \\
=& \frac{1}{V_{\mathrm{BZ}}^2} \sum_{m,l} \int_{\mathrm{BZ}} \mathrm{d}^3k \int_{\mathrm{BZ}} \mathrm{d}^3k' \, \mathcal{U}_{ln}^{(\mathbf{k}')} \, \langle\, \mathbf{H}_{l\mathbf{k}'}{}^{*} \,|\, \mathbf{H}_{m\mathbf{k}} \,\rangle \, \mathcal{U}_{mn}^{(\mathbf{k})} \\
=& \frac{1}{V_{\mathrm{BZ}}^2} \sum_{m,l} \int_{\mathrm{BZ}} \mathrm{d}^3k \int_{\mathrm{BZ}} \mathrm{d}^3k' \, \mathcal{U}_{ln}^{(\mathbf{k}')} \, \langle\, \mathbf{H}_{l(-\mathbf{k}')} \,|\, \mathbf{H}_{m\mathbf{k}} \,\rangle \, \mathcal{U}_{mn}^{(\mathbf{k})} \\
=& \frac{1}{V_{\mathrm{BZ}}} \sum_{m} \int_{\mathrm{BZ}} \mathrm{d}^3k \, \mathcal{U}_{mn}^{(-\mathbf{k})} \, \mathcal{U}_{mn}^{(\mathbf{k})}.
\end{aligned} \tag{A.6}$$

Here, we assumed time reversal symmetry to perform the step $|\, \mathbf{H}_{l\mathbf{k}}^{*} \,\rangle = |\, \mathbf{H}_{l(-\mathbf{k})} \,\rangle$. Finally, we find an expression for the norm of the imaginary parts:

$$\Omega_{\mathrm{imag}}^{(n)} = \frac{1}{2} \left[1 - \mathrm{Re} \left\{ \frac{1}{V_{\mathrm{BZ}}} \sum_{m} \int_{\mathrm{BZ}} \mathrm{d}^3k \, \mathcal{U}_{mn}^{(-\mathbf{k})} \, \mathcal{U}_{mn}^{(\mathbf{k})} \right\} \right]. \tag{A.7}$$

For a discretized **k**-mesh this becomes:

$$\Omega_{\mathrm{imag}}^{(n)} = \frac{1}{2} \left[1 - \frac{1}{N_k} \sum_{\mathbf{k}m} \mathrm{Re} \left\{ \mathcal{U}_{mn}^{(-\mathbf{k})} \mathcal{U}_{mn}^{(\mathbf{k})} \right\} \right]. \tag{A.8}$$

The Penalty Gradient

The norm of the imaginary parts of the band group reads:

$$\begin{aligned}
\Omega_{\mathrm{imag}} &= \sum_{n} \Omega_{\mathrm{imag}}^{(n)} && \text{(A.9)} \\
&= \frac{1}{2} \left[1 - \frac{1}{V_{\mathrm{BZ}}} \int_{\mathrm{BZ}} \mathrm{d}^3k \; \mathrm{Re} \Big\{ \sum_{mn} \mathcal{U}_{mn}^{(-\mathbf{k})} \mathcal{U}_{mn}^{(\mathbf{k})} \Big\} \right] && \text{(A.10)} \\
&= \frac{1}{2} \left[1 - \frac{1}{V_{\mathrm{BZ}}} \int_{\mathrm{BZ}} \mathrm{d}^3k \; \mathrm{Re}\, \mathrm{Tr} [\mathcal{U}^{(-\mathbf{k})} \cdot (\mathcal{U}^{(\mathbf{k})})^T] \right]. && \text{(A.11)}
\end{aligned}$$

Under an infinitesimal update of the mixing matrices

$$\mathcal{U}^{(\mathbf{k})} \mapsto \mathcal{U}^{(\mathbf{k})} (1 + \mathrm{d}\mathcal{W}^{(\mathbf{k})}) = \mathcal{U}^{(\mathbf{k})} + \mathcal{U}^{(\mathbf{k})} \cdot \mathrm{d}\mathcal{W}^{(\mathbf{k})}, \tag{A.12}$$

the imaginary part norm changes according to:

$$\begin{aligned}\Omega_{\mathrm{imag}} \mapsto\ & \frac{1}{2}\Big[1 - \frac{1}{V_{\mathrm{BZ}}}\int_{\mathrm{BZ}} \mathrm{d}^3k \ \mathrm{Re}\ \mathrm{Tr}\Big[(\mathcal{U}^{(-\mathbf{k})} + \mathcal{U}^{(-\mathbf{k})} \cdot \mathrm{d}\mathcal{W}^{(-\mathbf{k})}) \\ & \quad \times (\mathcal{U}^{(\mathbf{k})} + \mathcal{U}^{(\mathbf{k})} \cdot \mathrm{d}\mathcal{W}^{(\mathbf{k})})^T\Big]\Big] \qquad \text{(A.13)} \\ =\ & \frac{1}{2}\Big[1 - \frac{1}{V_{\mathrm{BZ}}}\int_{\mathrm{BZ}} \mathrm{d}^3k \ \mathrm{Re}\ \mathrm{Tr}\Big[\Big(\mathcal{U}^{(-\mathbf{k})} + \mathcal{U}^{(-\mathbf{k})} \cdot \mathrm{d}\mathcal{W}^{(-\mathbf{k})}\Big) \\ & \quad \times \Big((\mathcal{U}^{(\mathbf{k})})^T + (\mathrm{d}\mathcal{W}^{(\mathbf{k})})^T \cdot (\mathcal{U}^{(\mathbf{k})})^T\Big)\Big]\Big] \qquad \text{(A.14)} \\ =\ & \frac{1}{2}\Big[1 - \frac{1}{V_{\mathrm{BZ}}}\int_{\mathrm{BZ}} \mathrm{d}^3k \ \mathrm{Re}\ \mathrm{Tr}\Big[\mathcal{U}^{(-\mathbf{k})} \cdot (\mathcal{U}^{(\mathbf{k})})^T + \mathcal{U}^{(-\mathbf{k})} \cdot \mathrm{d}\mathcal{W}^{(-\mathbf{k})} \cdot (\mathcal{U}^{(\mathbf{k})})^T \\ & \quad + \mathcal{U}^{(-\mathbf{k})} \cdot (\mathrm{d}\mathcal{W}^{(\mathbf{k})})^T \cdot (\mathcal{U}^{(\mathbf{k})})^T + \mathcal{O}(\mathrm{d}\mathcal{W}^2)\Big]\Big]. \qquad \text{(A.15)}\end{aligned}$$

Therefore, the infinitesimal change reads:

$$\begin{aligned}\mathrm{d}\Omega_{\mathrm{imag}} =\ & -\frac{1}{2V_{\mathrm{BZ}}}\int_{\mathrm{BZ}} \mathrm{d}^3k \ \Big\{\mathrm{Re}\ \mathrm{Tr}[\mathcal{U}^{(-\mathbf{k})} \cdot \mathrm{d}\mathcal{W}^{(-\mathbf{k})} \cdot (\mathcal{U}^{(\mathbf{k})})^T] + \\ & \mathrm{Re}\ \mathrm{Tr}[\mathcal{U}^{\mathbf{k}} \cdot \mathrm{d}\mathcal{W}^{(\mathbf{k})} \cdot (\mathcal{U}^{(-\mathbf{k})})^T]\Big\} \qquad \text{(A.16)} \\ =\ & -\frac{1}{2V_{\mathrm{BZ}}}\int_{\mathrm{BZ}} \mathrm{d}^3k \ \Big\{\mathrm{Re}\ \mathrm{Tr}[\mathrm{d}\mathcal{W}^{(-\mathbf{k})} \cdot (\mathcal{U}^{(\mathbf{k})})^T \cdot \mathcal{U}^{(-\mathbf{k})}] + \\ & \mathrm{Re}\ \mathrm{Tr}[\mathrm{d}\mathcal{W}^{(\mathbf{k})} \cdot (\mathcal{U}^{(-\mathbf{k})})^T \cdot \mathcal{U}^{(\mathbf{k})}]\Big\} \qquad \text{(A.17)} \\ =\ & -\frac{1}{V_{\mathrm{BZ}}}\int_{\mathrm{BZ}} \mathrm{d}^3k \ \mathrm{Re}\ \mathrm{Tr}[\mathrm{d}\mathcal{W}^{(\mathbf{k})} \cdot (\mathcal{U}^{(-\mathbf{k})})^T \cdot \mathcal{U}^{(\mathbf{k})}]. \qquad \text{(A.18)}\end{aligned}$$

From the identity

$$\frac{\mathrm{d}\ \mathrm{Re}\,\mathrm{Tr}[\mathrm{d}\mathcal{W} \cdot \mathcal{B}]}{\mathrm{d}\mathcal{W}} = \mathcal{A}[\mathcal{B}] = \frac{1}{2}[\mathcal{B} - \mathcal{B}^\dagger] \qquad \text{(A.19)}$$

we eventually obtain the final expression for the gradient:

$$\frac{\mathrm{d}\Omega_{\mathrm{imag}}}{\mathrm{d}\mathcal{W}^{(\mathbf{k})}} = -\frac{1}{V_{\mathrm{BZ}}}\mathcal{A}[(\mathcal{U}^{(-\mathbf{k})})^T \cdot \mathcal{U}^{(\mathbf{k})}]. \qquad \text{(A.20)}$$

A.2. Point Symmetries

First, I would like to provide some intermediate steps for deriving Eqn. (3.51). We begin with Eqn. (3.50):

$$\int_{\mathrm{BZ}} \mathrm{d}k \sum_{ml} |\, H_{l\mathbf{k}}\,\rangle \mathcal{U}^{(\mathbf{k})}_{lm} \mathcal{D}^{(\hat{T},\mathrm{WF})}_{mn} = \int_{\mathrm{BZ}} \mathrm{d}k \sum_{ml} |\, H_{l\mathbf{k}}\,\rangle \mathcal{D}^{(\hat{T},\mathrm{BF},\hat{T}^{-1}\mathbf{k})}_{lm} \mathcal{U}^{(\hat{T}^{-1}\mathbf{k})}_{mn}, \qquad \text{(A.21)}$$

and project on some Bloch mode $|\, H_{n_0\mathbf{k}_0} \rangle$:

$$\begin{aligned} &\int_{\mathrm{BZ}} \mathrm{d}k \sum_{ml} \langle\, H_{n_0\mathbf{k}_0} \,|\, H_{l\mathbf{k}} \,\rangle \mathcal{U}^{(\mathbf{k})}_{lm} \mathcal{D}^{(\hat{T},\mathrm{WF})}_{mn} \\ &\quad = \int_{\mathrm{BZ}} \mathrm{d}k \sum_{ml} \langle\, H_{n_0\mathbf{k}_0} \,|\, H_{l\mathbf{k}} \,\rangle \mathcal{D}^{(\hat{T},\mathrm{BF},\hat{T}^{-1}\mathbf{k})}_{lm} \mathcal{U}^{(\hat{T}^{-1}\mathbf{k})}_{mn}. \end{aligned} \tag{A.22}$$

Because of the Bloch modes' orthonormality $\langle\, H_{n\mathbf{k}} \,|\, H_{n'\mathbf{k}'} \,\rangle = \delta_{nn'}\delta(\mathbf{k}-\mathbf{k}')$ we find:

$$\sum_m \mathcal{U}^{(\mathbf{k}_0)}_{n_0 m} \mathcal{D}^{(\hat{T},\mathrm{WF})}_{mn} = \sum_m \mathcal{D}^{(\hat{T},\mathrm{BF},\hat{T}^{-1}\mathbf{k}_0)}_{n_0 m} \mathcal{U}^{(\hat{T}^{-1}\mathbf{k}_0)}_{mn}, \tag{A.23}$$

which is equivalent to Eqn. (3.51).
Now we can derive the gradient for the symmetry–favoring penalty term Eqn. (3.52):

$$\Omega_{\mathrm{sym}} = \sum_{\hat{T}\in\mathcal{P}} \int_{\mathrm{BZ}} \mathrm{d}k \left(N_b - \mathrm{Tr}\left\{ [\mathcal{U}^{(\hat{T}^{-1}\mathbf{k})}]^{-1} [\mathcal{D}^{(\hat{T},\mathrm{BF},\hat{T}^{-1}\mathbf{k})}]^{-1} \mathcal{U}^{(\mathbf{k})} \mathcal{D}^{(\hat{T},\mathrm{WF})} \right\} \right). \tag{A.24}$$

As before we apply an infinitesimal change of the unitary mixing:

$$\mathcal{U}^{\mathbf{k}} \mapsto \quad \mathcal{U}^{\mathbf{k}} \cdot (1 + \mathrm{d}\mathcal{W}^{\mathbf{k}}) \quad = \mathcal{U}^{\mathbf{k}} + \mathcal{U}^{\mathbf{k}} \cdot \mathrm{d}\mathcal{W}^{\mathbf{k}}, \tag{A.25}$$

$$[\mathcal{U}^{\mathbf{k}}]^{-1} \mapsto \quad (1 - \mathrm{d}\mathcal{W}^{\mathbf{k}}) \cdot [\mathcal{U}^{\mathbf{k}}]^{-1} \quad = [\mathcal{U}^{\mathbf{k}}]^{-1} - \mathrm{d}\mathcal{W}^{\mathbf{k}} \cdot [\mathcal{U}^{\mathbf{k}}]^{-1}. \tag{A.26}$$

This causes a change of the penalty term:

$$\begin{aligned} \Omega_{\mathrm{sym}} \quad \mapsto \quad & \sum_{\hat{T}\in\mathcal{P}} \int_{\mathrm{BZ}} \mathrm{d}k \Big(N_b - \mathrm{Tr}\Big\{ (1 - \mathrm{d}\mathcal{W}^{(\hat{T}^{-1}\mathbf{k})})\ [\mathcal{U}^{(\hat{T}^{-1}\mathbf{k})}]^{-1} \\ & \qquad \times [\mathcal{D}^{(\hat{T},\mathrm{BF},\hat{T}^{-1}\mathbf{k})}]^{-1} \mathcal{U}^{(\mathbf{k})} (1 + \mathrm{d}\mathcal{W}^{(\mathbf{k})}) \mathcal{D}^{(\hat{T},\mathrm{WF})} \Big\} \Big), \end{aligned} \tag{A.27}$$

$$\begin{aligned} = \quad & \sum_{\hat{T}\in\mathcal{P}} \int_{\mathrm{BZ}} \mathrm{d}k \Big(N_b - \mathrm{Tr}\Big\{ [\mathcal{U}^{(\hat{T}^{-1}\mathbf{k})}]^{-1} [\mathcal{D}^{(\hat{T},\mathrm{BF},\hat{T}^{-1}\mathbf{k})}]^{-1} \mathcal{U}^{(\mathbf{k})} \mathcal{D}^{(\hat{T},\mathrm{WF})} \Big\} \\ & \qquad + \mathrm{Tr}\Big\{ \mathrm{d}\mathcal{W}^{(\hat{T}^{-1}\mathbf{k})} [\mathcal{U}^{(\hat{T}^{-1}\mathbf{k})}]^{-1} [\mathcal{D}^{(\hat{T},\mathrm{BF},\hat{T}^{-1}\mathbf{k})}]^{-1} \mathcal{U}^{(\mathbf{k})} \mathcal{D}^{(\hat{T},\mathrm{WF})} \Big\} \\ & \qquad - \mathrm{Tr}\Big\{ \mathrm{d}\mathcal{W}^{(\mathbf{k})} \mathcal{D}^{(\hat{T},\mathrm{WF})} [\mathcal{U}^{(\hat{T}^{-1}\mathbf{k})}]^{-1} [\mathcal{D}^{(\hat{T},\mathrm{BF},\hat{T}^{-1}\mathbf{k})}]^{-1} \mathcal{U}^{(\mathbf{k})} \Big\} \\ & \qquad + \mathcal{O}(\mathrm{d}\mathcal{W}^2) \Big). \end{aligned} \tag{A.28}$$

By collecting all terms linear in $\mathrm{d}\mathcal{W}^{(\mathbf{k})}$ for a specific wave vector $\mathbf{k}$, and using the identity

$$\frac{\mathrm{dTr}[\mathrm{d}\mathcal{W} \cdot \mathcal{B}]}{\mathrm{d}\mathcal{W}} = \mathcal{B} \tag{A.29}$$

we finally arrive at the desired gradient:

$$\begin{aligned} \frac{\mathrm{d}\Omega_{\mathrm{sym}}}{\mathrm{d}\mathcal{W}^{(\mathbf{k})}} \quad = \quad & [\mathcal{U}^{(\mathbf{k})}]^{-1} [\mathcal{D}^{(\hat{T},\mathrm{BF},\mathbf{k})}]^{-1} \mathcal{U}^{(\hat{T}\mathbf{k})} \mathcal{D}^{(\hat{T},\mathrm{WF})} \\ & - \mathcal{D}^{(\hat{T},\mathrm{WF})} [\mathcal{U}^{(\hat{T}^{-1}\mathbf{k})}]^{-1} [\mathcal{D}^{(\hat{T},\mathrm{BF},\hat{T}^{-1}\mathbf{k})}]^{-1} \mathcal{U}^{(\mathbf{k})}. \end{aligned} \tag{A.30}$$

B

Appendix B.

Numerical Lamb Shift Integration

For numerically evaluating the Lamb shift integral Eqn. (5.31), we linearly approximate the projected local density of states on an interval (ω_i, ω_{i+1}):

$$N_l^{(\mathrm{num})}(\omega) \approx N_i^{(0)} + N_i^{(1)}\omega. \tag{B.1}$$

If the pole is not inside this integration domain, i. e. , $\omega_0 \notin (\omega_i, \omega_{i+1})$, the principal value simply becomes an ordinary integral that can be easily evaluated:

$$\begin{aligned}
&\int_{\omega_i}^{\omega_{i+1}} \mathrm{d}\omega' \frac{\gamma(\omega')}{\omega_0 - \omega'} \\
=&\int_{\omega_i}^{\omega_{i+1}} \frac{\mathrm{d}\omega'}{\omega_0 - \omega'} \left(N_i^{(0)}\omega' + N_i^{(1)}\omega'^2 \right) && \text{(B.2)} \\
=&N_i^{(0)} \int_{\omega_i}^{\omega_{i+1}} \mathrm{d}\omega' \left[\frac{\omega_0}{\omega_0 - \omega'} - 1 \right] \\
&\quad + N_i^{(1)} \int_{\omega_i}^{\omega_{i+1}} \mathrm{d}\omega' \left[\frac{\omega_0^2}{\omega_0 - \omega'} - \omega_0 - \omega' \right] && \text{(B.3)} \\
=&\left(N_i^{(0)}\omega_0 + N_i^{(1)}\omega_0^2 \right) \int_{\omega_i}^{\omega_{i+1}} \frac{\mathrm{d}\omega'}{\omega_0 - \omega'} \\
&\quad - \left(N_i^{(0)} + N_i^{(1)}\omega_0 \right) \int_{\omega_i}^{\omega_{i+1}} \mathrm{d}\omega' && \text{(B.4)} \\
&\quad - N_i^{(1)} \int_{\omega_i}^{\omega_{i+1}} \omega' \mathrm{d}\omega' \\
=&\left(N_i^{(0)}\omega_0 + N_i^{(1)}\omega_0^2 \right) \left(\log\left|\omega_0 - \omega_{i+1}\right| - \log\left|\omega_0 - \omega_i\right| \right) \\
&\quad + \left(N_i^{(0)} + N_i^{(1)}\omega_0 \right) \left(\omega_i - \omega_{i+1} \right) + \frac{N_i^{(1)}}{2} \left(\omega_i^2 - \omega_{i+1}^2 \right). && \text{(B.5)}
\end{aligned}$$

However, if the pole lies inside the integration domain, the principal value must be evaluated according to its definition as a limit of the sum of sub–integrals:

$$\begin{aligned}
&\mathfrak{p}\int_{\omega_n}^{\omega_{n+1}} \mathrm{d}\omega' \frac{\gamma(\omega')}{\omega_0-\omega'} \\
&= \lim_{h\to 0}\left[\int_{\omega_n}^{\omega_0-h} \mathrm{d}\omega' \frac{\gamma(\omega')}{\omega_0-\omega'} + \int_{\omega_0+h}^{\omega_{n+1}} \mathrm{d}\omega' \frac{\gamma(\omega')}{\omega_0-\omega'}\right] && \text{(B.6)}\\
&=\left(N_i^{(0)}\omega_0 + N_i^{(1)}\omega_0^2\right) \\
&\qquad\times \lim_{h\to 0}\left(\log\left|\omega_0-\omega_{n+1}\right| - \log\left|-h\right| + \log\left|h\right| - \log\left|\omega_0-\omega_n\right|\right) \\
&\qquad+ \left(N_i^{(0)} + N_i^{(1)}\omega_0\right)\lim_{h\to 0}\left(\omega_i - \omega_0 + h + \omega_0 + h - \omega_{i+1}\right) && \text{(B.7)}\\
&\qquad+ \frac{N_i^{(1)}}{2}\lim_{h\to 0}\left(\omega_i^2 - \omega_0^2 + 2h\omega_0 - h^2 + \omega_0^2 + 2h\omega_0 + h^2 - \omega_{i+1}^2\right) \\
&=\left(N_i^{(0)}\omega_0 + N_i^{(1)}\omega_0^2\right)\left(\log\left|\omega_0-\omega_{i+1}\right| - \log\left|\omega_0-\omega_i\right|\right) \\
&\qquad+ \left(N_i^{(0)} + N_i^{(1)}\omega_0\right)\left(\omega_i - \omega_{i+1}\right) + \frac{N_i^{(1)}}{2}\left(\omega_i^2 - \omega_{i+1}^2\right) && \text{(B.8)}
\end{aligned}$$

Conveniently, this is exactly the same expression as Eqn. (B.5). The distinction whether ω_0 is greater or smaller than any of the integration boundaries is handled by the absolute value. However, obviously, the result ceases to be defined if $\omega_0 = \omega_{i/i+1}$.

Bibliography

[1] E. Yablonovitch, “Inhibited spontaneous emission in solid–state physics and electronics,” *Phys. Rev. Lett.*, vol. 58, pp. 2059–2062, 1987.

[2] S. John, “Strong localization of photons in certain disordered dielectric superlattices,” *Phys. Rev. Lett.*, vol. 58, pp. 2486–2489, 1987.

[3] E. Cubukcu, K. Aydin, E. Özbay, S. Foteinopoulou, and C. M. Soukoulis, “Subwavelength Resolution in a Two–Dimensional Photonic–Crystal–Based Superlens,” *Phys. Rev. Lett.*, vol. 91, p. 207401, 2003.

[4] H. Kosaka, T. Kawashima, A. Tomita, M. Notomi, T. Tamamura, T. Sato, and S. Kawakami, “Superprism phenomena in photonic crystals,” *Phys. Rev. B*, vol. 58, pp. R10096–R10099, 1998.

[5] S. Y. Lin, J. Moreno, and J. G. Fleming, “Three–dimensional photonic–crystal emitter for thermal photovoltaic power generation,” *Appl. Phys. Lett.*, vol. 83, p. 380, 2003.

[6] C. Schuler, “Thermal Emission of Photonic Crystals,” Master’s thesis, Karlsruhe Institute of Technology, 2009.

[7] M. Lončar, T. Yoshie, A. Scherer, P. Gogna, and Y. Qiu, “Low–threshold photonic crystal laser,” *Appl. Phys. Lett.*, vol. 81, p. 2680, 2002.

[8] H. Nakamura, Y. Sugimoto, and K. Asakawa, “Ultra–fast photonic crystal/quantum dot all–optical switch for future photonic networks,” in *Conference on Lasers and Electro–Optics/Quantum Electronics and Laser Science Conference and Photonic Applications Systems Technologies*, p. CThK1, Optical Society of America, 2006.

[9] A. Taflove and S. C. Hagness, *Computational Electrodynamics: The Finite–Difference Time–Domain Method.* Artech House, Boston, 3. ed. ed., 2005.

[10] K. Busch, M. König, and J. Niegemann, “Discontinuous Galerkin methods in nanophotonics,” *Laser and Photonics Reviews*, vol. 5, p. 773–809, 2011.

[11] L. Li, “Use of Fourier series in the analysis of discontinuous periodic structures,” *J. Opt. Soc. Am. A*, vol. 13, pp. 1870–1876, 1996.

[12] S. Essig and K. Busch, “Generation of adaptive coordinates and their use in the Fourier Modal Method,” *Opt. Express*, vol. 18, p. 23258, 2010.

[13] K. Busch, C. Blum, A. M. Graham, D. Hermann, M. Köhl, P. Mack, and C. Wolff, "The photonic Wannier function approach to photonic crystal simulations: status and perspectives," *J. Mod. Optics*, vol. 58, p. 365, 2011.

[14] M. Köhl, "Ungeordnete photonische Kristalle: Eine Untersuchung basierend auf Green'schen Funktionen und der Wannierbasis," Master's thesis, Karlsruhe Institute of Technology, 2010.

[15] M. Köhl, C. Wolff, and K. Busch, "Cluster Coherent Potential Approximation for Disordered Photonic Crystals using Photonic Wannier Functions," *Opt. Lett.*, (submitted).

[16] S. G. Johnson and J. D. Joannopoulos, "Block–iterative frequency-domain methods for Maxwell's equations in a planewave basis," *Opt. Express*, vol. 8, pp. 173–190, 2001.

[17] S. John and T. Quang, "Spontaneous emission near the edge of a photonic band gap," *Phys. Rev. A*, vol. 50, pp. 1764–1769, 1994.

[18] Y. Akahane, T. Asano, B.-S. Song, and S. Noda, "High–Q photonic nanocavity in a two-dimensional photonic crystal," *Nature*, vol. 425, pp. 944–947, 2003.

[19] J. D. Jackson, *Classical electrodynamics.* Wiley, 3. ed. ed., 1999.

[20] M. Lax, *Symmetry Principles in Solid State and Molecular Physics.* New York, NY, USA: Dover Publications, Inc., 1974.

[21] K. Sakoda, *Optical Properties of Photonic Crystals.* Springer, 2001.

[22] P. Mack, *2D H–polarized Auxiliary Basis Functions for the Extension of the Photonic Wannier Function Expansion for Photonic Crystal Circuitry.* PhD thesis, Karlsruhe Institute of Technology, 2011.

[23] N. W. Ashcroft and N. D. Mermin, *Solid state physics.* Brooks/Coole Thomson Learning, 2005.

[24] K. M. Ho, C. T. Chan, C. M. Soukoulis, R. Biswas, and M. Sigalas, "Photonic band gaps in three dimensions New layer–by–layer periodic structures," *Solid State Communications*, vol. 89, pp. 413–416, 1994.

[25] C. Schuler, C. Wolff, K. Busch, and M. Florescu, "Thermal emission from finite photonic crystals," *Appl. Phys. Lett.*, vol. 95, p. 241103, 2009.

[26] T. Yoshie, A. Scherer, J. Hendrickson, G. Khitrova, H. M. Gibbs, G. Rupper, C. Ell, O. B. Shchekin, and D. G. Deppe, "Vacuum Rabi splitting with a single quantum dot in a photonic crystal nanocavity," *Nature*, vol. 432, pp. 200–203, 2004.

[27] M. Deubel, M. Wegener, S. Linden, G. von Freymann, and S. John, "3D–2D–3D photonic crystal heterostructures fabricated by direct laser writing," *Opt. Lett.*, vol. 31, pp. 805–807, 2006.

[28] F. Intonti, S. Vignolini, V. Türck, M. Colocci, P. Bettotti, L. Pavesi, S. L. Schweizer, R. Wehrspohn, and D. Wiersma, "Rewritable photonic circuits," *Appl. Phys. Lett.*, vol. 89, p. 211117, 2006.

[29] M. H. Schillinger, *Maximally localized photonic Wannier functions for the highly efficient description of intgrated Photonic Circuits.* PhD thesis, Universität Karlsruhe, 2006.

[30] D. Hermann, *Wannier–Function Based Scattering–Matrix Formalism for Photonic Crystal Circuitry.* PhD thesis, Universität Karlsruhe, 2008.

[31] D. Hermann, M. Diem, S. F. Mingaleev, A. Garcia-Martin, P. Wölfle, and K. Busch, "Photonic crystals with anomalous dispersion: Unconventional propagating modes in the photonic band gap," *Phys. Rev. B*, vol. 77, p. 035112, 2008.

[32] K. Busch and S. John, "Photonic band gap formation in certain self–organizing systems," *Phys. Rev. E*, vol. 58, pp. 3896–3908, 1998.

[33] J. M. Harrison, P. Kuchment, A. Sobolev, and B. Winn, "On occurrence of spectral edges for periodic operators inside the Brillouin zone," *J. Phys. A: Math. Theor.*, vol. 40, p. 7597, 2007.

[34] G. Lehmann and M. Taut, "On the Numerical Calculation of the Density of States and Related Properties," *phys. stat. sol. (b)*, vol. 54, pp. 469–477, 1972.

[35] J. Hama, M. Watanabe, and T. Kato, "Correctly weighted tetrahedron method for k–space integration ," *J. Phys. Condens. Matter*, vol. 2, pp. 7445–7452, 1990.

[36] H. J. Monkhorst and J. D. Pack, "Special points for Brillouin–zone integrations," *Phys. Rev. B*, vol. 13, pp. 5188–5192, 1976.

[37] N. Marzari and D. Vanderbilt, "Maximally localized generalized Wannier functions for composite energy bands," *Phys. Rev. B*, vol. 56, pp. 12847–12865, 1997.

[38] I. Souza, N. Marzari, and D. Vanderbilt, "Maximally localized Wannier functions for entangled energy bands," *Phys. Rev. B*, vol. 65, p. 035109, 2001.

[39] A. Garcia-Martin, D. Hermann, K. Busch, and P. Wölfle, "Defect computation in photonic crystals: a solid state theoretical approach," *Nanotechnology*, vol. 14, p. 177, 2003.

[40] K. Busch, S. F. Mingaleev, M. Schillinger, and D. Hermann, "The Wannier function approach to photonic crystal circuits," *J. Phys. Condens. Matter*, vol. 15, p. R1233, 2003.

[41] H. Takeda, A. Chutinan, and S. John, "Localized light orbitals: Basis states for three–dimensional photonic crystal microscale circuits," *Phys. Rev. B*, vol. 74, p. 195116, 2006.

[42] J. E. Sipe, "Vector $k \cdot p$ approach for photonic band structures," *Phys. Rev. E*, vol. 62, pp. 5672–5677, 2000.

[43] C. Geiger and C. Kanzow, *Numerische Verfahren zur Lösung unrestringierter Optimierungsaufgaben*. Springer, 1999.

[44] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by Simulated Annealing," *Science*, vol. 220, pp. 671–680, 1983.

[45] C. Wolff and K. Busch, "Constraints in the Generation of Photonic Wannier Functions," *Physica B*, (in press).

[46] M. Eisenberg and R. Guy, "A Proof of the Hairy Ball Theorem," *Am. Math. Mon.*, vol. 86, p. 571–574, 1979.

[47] I. Staude, G. von Freymann, S. Essig, K. Busch, and M. Wegener, "Waveguides in three–dimensional photonic–bandgap materials by direct laser writing and silicon double inversion," *Opt. Lett.*, vol. 36, p. 67, 2011.

[48] C. Blum, "Time–Domain Method for the Solution of Maxwell's Equations in Periodic Media using Wannier Functions," Master's thesis, Karlsruhe Institute of Technology, 2010.

[49] C. Blum, C. Wolff, and K. Busch, "Photonic–crystal time–domain simulations using Wannier functions," *Opt. Lett.*, vol. 36, p. 307, 2011.

[50] A. M. Graham, "Time Domain Simulations of Non–Linear Photonic Crystals using Wannier Functions," Master's thesis, Karlsruhe Institute of Technology, 2010.

[51] A. F. Oskooi, L. Zhang, Y. Avniel, and S. G. Johnson, "The failure of perfectly matched layers, and towards their redemptionby adiabatic absorbers," *Opt. Express*, vol. 16, pp. 11376–11392, 2008.

[52] J. F. Galisteo-López, E. Palacios-Lidón, E. Castillo-Martínez, and C. López, "Optical study of the pseudogap in thickness and orientation controlled artificial opals," *Phys. Rev. B*, vol. 68, p. 115109, 2003.

[53] J. F. Galisteo-López, M. Galli, M. Patrini, A. Balestreri, L. C. Andreani, and C. López, "Effective refractive index and group velocity determination of three–dimensional photonic crystals by means of white light interferometry," *Phys. Rev. B*, vol. 73, p. 125103, 2006.

[54] S. Romanov, U. Peschel, M. Bardosova, S. Essig, and K. Busch, “Suppression of the critical angle of diffraction in thin–film colloidal photonic crystals,” *Phys. Rev. B*, vol. 82, p. 115403, 2010.

[55] F. Genereux, S. W. Leonard, H. M. van Driel, A. Birner, and U. Gösele, “Large birefringence in two–dimensional silicon photonic crystals,” *Phys. Rev. B*, vol. 63, p. 161101, 2001.

[56] A. Reza, Z. Balevicius, R. Vaisnoras, G. Babonas, and A. Ramanavicius, “Studies of optical anisotropy in opals by normal incidence ellipsometry,” *Thin Solid Films*, vol. 519, pp. 2641–2644, 2011. 5th International Conference on Spectroscopic Ellipsometry (ICSE-V).

[57] M. Born and E. Wolf, *Principles of Optics*. Pergamon Press, 1959.

[58] A. Balestreri, L. C. Andreani, and M. Agio, “Optical properties and diffraction effects in opal photonic crystals,” *Phys. Rev. E*, vol. 74, p. 036603, 2006.

[59] C. Wolff, S. G. Romanov, J. Küchenmeister, U. Peschel, and K. Busch, “Polarization Dynamics in Cubic Photonic Crystals,” *Phys. Rev. Lett.*, (submitted).

[60] X. Checoury, S. Enoch, C. Lopez, and A. Blanco, “Stacking patterns in self–assembly opal photonic crystals,” *Appl. Phys. Lett.*, vol. 90, p. 161131, 2007.

[61] U. Hoeppe, C. Wolff, J. Küchenmeister, J. Niegemann, M. Drescher, H. Benner, and K. Busch, “Direct observation of non–Markovian radiation dynamics in 3D bulk photonic crystals,” *Phys. Rev. Lett.*, vol. 108, p. 043603, 2012.

[62] S. John and J. Wang, “Quantum electrodynamics near a photonic band gap Photon bound states and dressed atoms,” *Phys. Rev. Lett.*, vol. 64, pp. 2418–2421, 1990.

[63] L. Sapienza, H. Thyrrestrup, S. Stobbe, P. Garcia, S. Smolka, and P. Lodahl, “Cavity Quantum Electrodynamics with Anderson–Localized Modes,” *Science*, vol. 327, pp. 1352–1355, 2010.

[64] I. S. Nikolaev, P. Lodahl, A. F. van Driel, A. F. Koenderink, and W. L. Vos, “Strongly nonexponential time–resolved fluorescence of quantum–dot ensembles in three–dimensional photonic crystals,” *Phys. Rev. B*, vol. 75, p. 115302, 2007.

[65] P. Lodahl, A. van Driel, I. Nikolaev, A. Irman, K. Overgaag, D. Vanmaekelbergh, and W. Vos, “Controlling the dynamics of spontaneous emission from quantum dots by photonic crystals,” *Nature*, vol. 430, pp. 654–657, 2004.

[66] Q. Liu, H. Song, W. Wang, X. Bai, Y. Wang, B. Dong, L. Xu, and W. Han, “Observation of Lamb shift and modified spontaneous emission dynamics in the YBO3Eu3+ inverse opal,” *Opt. Lett.*, vol. 35, p. 2898, 2010.

[67] W. L. Hahn, “Concepts of NMR in quantum optics,” *Concepts in Magnetic Resonance*, vol. 9, pp. 69–81, 1997.

[68] E. Özbay, A. Abeyta, G. Tuttle, M. Tringides, W. Biswas, C. T. Chan, C. M. Soukoulis, and M. M. Ho, “Measurement of a three–dimensional photonic band gap in a crystal structure made of dielectric rods,” *Phys. Rev. B*, vol. 50, pp. 1945–1948, 1994.

[69] K. Busch, N. Vats, S. John, and B. C. Sanders, “Radiating dipoles in Photonic Crystals,” *Phys. Rev. E*, vol. 62, p. 4251, 2000.

[70] N. Vats, S. John, and K. Busch, “Theory of fluorescence in photonic crystals,” *Phys. Rev. A*, vol. 65, p. 043808, 2002.

[71] L. Mandel and E. Wolf, *Optical coherence and quantum optics.* Cambridge University Press, 1995.

[72] W. H. L., *Quantum Statistical Properties of Radiation.* John Wiley & Sons Inc, 1973.

[73] H. Thyrrestrup, L. Sapienza, and P. Lodahl, “Extraction of the β–factor for single quantum dots coupled to a photonic crystal waveguide,” *Appl. Phys. Lett.*, vol. 96, p. 231106, 2010.

[74] J. E. Mercereau, “Ferromagnetic Resonance g Factor to Order $(kRo)^2$,” *J. Appl. Phys.*, vol. 30, p. 184S, 1959.

[75] V. H. Wigner and E. P. Weisskopf *Z. Phys.*, vol. 63, p. 54, 1930.

[76] A. Z. Genack and N. Garcia, “Observation of photon localization in a three–dimensional disordered system,” *Phys. Rev. Lett.*, vol. 66, pp. 2064–2067, 1991.

Publications

In Peer–Reviewed Journals

- C. Wolff and K. Busch, “Constraints in the Generation of Photonic Wannier Functions”, *Physica B: solid state physics*, (in press).
- C. Wolff, S. G. Romanov, J. Küchenmeister, U. Peschel and K. Busch, “Polarization Dynamics in Cubic Photonic Crystals”, *Phys. Rev. Lett.*, (submitted).
- U. Hoeppe, C. Wolff, J. Küchenmeister, J. Niegemann, M. Drescher, H. Benner and K. Busch, “Direct Observation of Non-Markovian Radiation Dynamics in 3D Bulk Photonic Crystals”, *Phys. Rev. Lett.*, vol. 108, p. 043603, 2012.
- K. Busch, C. Blum, A. M. Graham, D. Hermann, M. Köhl, P. Mack and C. Wolff, “The photonic Wannier function approach to photonic crystal simulations: status and perspectives”, *J. Mod. Optics*, vol. 58, p. 365, 2011.
- C. Blum, C. Wolff and K. Busch, “Photonic-Crystal Time-Domain Simulations”, *Opt. Lett.*, vol. 36, p. 307, 2011.
- M. Köhl, C. Wolff and K. Busch, “Cluster Coherent Potential Approximation for Disordered Photonic Crystals using Photonic Wannier Functions”, *Opt. Lett.*, (submitted).
- I. Staude, M. Thiel, S. Essig, C. Wolff, K. Busch, G. von Freymann and M. Wegener, “Fabrication and characterization of silicon woodpile photonic crystals with a complete bandgap at telecom wavelengths”, *Opt. Lett.*, vol. 35, p. 1094, 2010.
- C. Schuler, C. Wolff, K. Busch and M. Florescu, “Thermal emission from finite photonic crystals”, *Appl. Phys. Lett.*, vol. 95, p. 241103, 2009.

Conferences and Workshops

- C. Wolff, S. G. Romanov, J. Küchenmeister, U. Peschel and K. Busch, “Polarization Change in Face–Centered Cubic Opal Films” (oral presentation), *TaCoNaPhotonics*, Bad Honnef, 2011;
  Proceedings Paper:
  C. Wolff, S. G. Romanov, J. Küchenmeister, U. Peschel and K. Busch, “Polarization Change in Face–Centered Cubic Opal Films”, *AIP Conference Proceedings*, vol. 1398, pp. 112–114, 2011.

- C. Wolff, J. Küchenmeister, K. Busch, S. G. Romanov and U. Peschel, "Cross–polarization breakdown in 3–dimensional opal photonic crystals" (oral presentation), *CLEO Europe*, Munich, 2011;
  Proceedings Paper:
  C. Wolff, J. Küchenmeister, K. Busch, S. G. Romanov and U. Peschel, "Cross–polarization breakdown in 3–dimensional opal photonic crystals",
  CLEO/Europe and EQEC 2011 Conference Digest CK9_6, 2011.

- C. Blum, C. Wolff and K. Busch, "The Wannier Function Approach to the Simulation of Photonic Crystal Defect Structures" (poster), *SPIE Europe*, Brussels, 2010.

- C. Wolff and K. Busch, "Generation of 3D Wannier Functions" (poster), *TaCoNa-Photonics*, Bad Honnef, 2009;
  Proceedings Paper:
  C. Wolff and K. Busch, "Generation of 3D Wannier Functions", *AIP Conference Proceedings*, vol. 1176, pp. 166–168, 2009.

- C. Wolff, P. Mack, C. Blum and K. Busch, "Wannier Functions for the Simulation of Photonic Crystal Structures" (oral presentation), 5th *Workshop on Numerical Methods for Optical Nano Structures*, Zurich, 2009.

- C. Wolff, C. Kölper, P. Mack, D. Hermann and K. Busch, "3D Vectorial Wannier Functions for the Design of Functional Defect Structures in Photonic Crystals" (oral presentation), *DPG-Frühjahrstagung*, Dresden, 2009.

- CFN Summer School on Nano–Photonics 2008.

- KSOP Summer School 2008.

Miscellaneous

- Idea, prototyping und essential contributions to the exhibit "Optische Schaltkreise mittels Photonischer Kristalle" in the exhibition "NANO! Nutzen und Visionen einer neuen Technologie" at the Technoseum Mannheim; furthermore:
  Article "Photonische Kristalle - Grundlage für Schaltkreise mit Licht" in the accompanying catalog.

- Development and fabrication of USB–compatible temperature sensors for monitoring the TFP/TKM server room and the physics department's pool rooms.

Acknowledgments

This thesis exists only partially due to my personal effort. It would have been impossible to write without the contributions and support of many people, most notably my family. Here, I would like to give those persons credit who were involved in my work:

- First, I would like to thank Prof. Kurt Busch who accepted my as a PhD student after my diploma, who introduced me into the field of photonic crystals, was always willing to answer any question I had and proved great patience with the progress of my work.
- I am also very grateful that Prof. Martin Wegener agreed to spend some of his precious time to co–referee my thesis.
- I am deeply indebted to Patrick Mack, Daniel Hermann and Matthias Schillinger for their great work in the context of the photonic Wannier method that I could base my own research on. Furthermore, I would like to thank them and Jens Niegemann for very inspiring discussions.
- The last two chapters of this thesis are the result of work performed together with experimentalists and other members of our group. For the close and pleasant collaborations, I would like to thank Jens Küchenmeister, Sergei Romanov, Ulf Peschel, Ulrich Hoeppe and Harmut Benner.
- This thesis would contain a vast number of typographic errors and incomprehensible expressions without the proof reading effort of Richard Diehl, Sabine Essig, Michael König, Jens Küchenmeister, Patrick Mack, Matthias Moeferdt, Lars Reichardt, Isabelle Wolff and Thomas Zebrowski.

Finally, I would like to express the great pleasure I had during the five years I spent in the group as a Hiwi, a diploma and a PhD student. Despite my personal quirks, I felt fully accepted in the group. I particularly enjoyed the regular barbecues (I like chicken salad!), and the vast amount of cake and ice cream we had during this time.

$$\Big\{(x,y) : 5 - 10(|x| - y)^2 - (|x| + y)^4 \geq 0\Big\}$$

`:wq`