



UNIVERSITÀ
DEGLI STUDI
FIRENZE

International Doctorate in Atomic and Molecular Photonics

Ciclo XXIX

Coordinatore Prof. Roberto Righini

**Imaging light transport at the femtosecond
scale: a walk on the wild side of diffusion**

Settore Scientifico Disciplinare FIS/03

Dottorando

Dott. Lorenzo Pattelli

Tutor

Prof. Diederik S. Wiersma

Coordinatore

Prof. Roberto Righini

2013–2016

Imaging light transport at the femtosecond scale: a walk on the wild side of diffusion

2016, Lorenzo Pattelli

International Doctorate in Atomic and Molecular Photonics, November 2016

Coordinator: Prof. Roberto Righini

Supervisor: Prof. Diederik S. Wiersma

Reviewers: Prof. Hui Cao and Prof. Frank Scheffold

Università degli Studi di Firenze

group of Optics of Complex Systems

Laboratorio Europeo di Spettroscopia non Lineare

Via Nello Carrara 1, Sesto Fiorentino

50019 - Florence, Italy

Contents

Acronyms	vii
Symbols	ix
1 Introduction	1
2 Light transport in turbid media	5
2.1 From Maxwell's equations to Radiative transfer	6
2.2 From Radiative transfer to Diffusion	22
3 Spatially-resolved time-of-flight spectroscopy	41
3.1 Optical characterization of turbid media	41
3.2 Cross-correlation optical gating	44
3.3 Ultrafast imaging	51
4 Experimental results	57
4.1 Validation of the technique	57
4.2 Unveiling data evaluation artifacts	61
4.3 A novel transport regime in ultra-thin samples	68
5 Space-time characterization of the ballistic-to-diffusive transition	75
5.1 Inverting light transport in a scattering slab	75
5.2 MCPLUSPLUS: a scriptable Monte Carlo library for radiative transfer . . .	77
5.3 Deconstructing light transport at the ballistic-to-diffusive transition . . .	81
5.4 Monte Carlo lookup-table based on spatio-temporal descriptors	88
6 Asymptotic transport in bounded media	93
6.1 Diffusive light transport in a semitransparent slab	93
6.2 Effective random-walk statistics	96
6.3 A walk on the wild side of diffusion	101
Bibliography	111
A Derivations	125
B Large-scale generation of exponentially distributed random numbers	133

List of publications

Publications related to this work

- Pattelli L., Savo R., Burresi M., & Wiersma D.S. Spatio-temporal visualization of light transport in complex photonic structures. *Light: Science & Applications* **5**, e16090 (2016)
- Mazzamuto G., Pattelli L., Toninelli C. & Wiersma D.S. Deducing effective light transport parameters in optically thin systems. *New Journal of Physics* **18**, 023036 (2016)
- Pattelli L., Mazzamuto G., Wiersma D.S. & Toninelli C. Diffusive light transport in semitransparent media. *Physical Review A* **94**, 043846 (2016)

Other publications

- Burresi M., Cortese L., Pattelli L., Kolle M., Vukusic P., Wiersma D.S., Steiner U. & Vignolini S. Bright-white beetle scales optimise multiple scattering of light. *Scientific Reports* **4** (2014)
- Cortese L., Pattelli L., Utel F., Vignolini S., Burresi M. & Wiersma D.S. Anisotropic light transport in white beetle scales. *Advanced Optical Materials* **3**, 1377–1341 (2015)
- Tiwari A.K., Boschetti A., Pattelli L., Zeng H., Torre R. & Wiersma D.S. Spectral super-resolution via stochastic mode competition in disordered media (2017, in preparation)
- Egel A., Pattelli L., Mazzamuto G., Wiersma D.S. & Lemmer U. CELES: CUDA-accelerated simulation of electromagnetic scattering by large ensembles of spheres. (2017, in preparation).

Acronyms

BBO	β -Barium borate.
CDF	cumulative distribution function.
CSR	central scaling region.
CTRW	continuous-time random walk.
CW	continuous wave.
DA	diffusion approximation.
DE	diffusion equation.
DPSS	diode-pumped solid-state.
EBC	extrapolated boundary condition.
EBPC	extrapolated boundary partial current.
HG	Heyney-Greenstein (phase function).
LUT	lookup table.
MC	Monte Carlo.
MSD	mean square displacement.
MSW	mean square width.
NIR	near infrared.
OOP	object-oriented programming.
OPA	optical parametric amplification.
OPO	optical parametric oscillator.
OT	optical thickness.
PCBC	partial current boundary condition.
PDF	probability density function.
PMT	photo-multiplier tube.
PRNG	pseudo-random number generator.
RTE	radiative transfer equation.
SEM	scanning electron microscope.
SFG	sum-frequency generation.
SHG	second-harmonic generation.
SLD	step-length distribution.
UFI	ultrafast imaging.

Symbols

ϖ_0	single particle albedo.
\mathbf{B}	magnetic field [T].
\mathbf{b}	unit vector along \mathbf{B} .
c	light speed in vacuum $c = \frac{1}{\sqrt{\epsilon_0 \mu_0}}$ [m s ⁻¹].
D	diffusion coefficient [m s ⁻²].
\mathbf{D}	diffusion coefficient [m s ⁻²].
$\Delta\tau_p$	full-width-half-maximum duration of a laser pulse [s].
\mathbf{E}	electric field [N C ⁻¹].
\mathbf{e}	unit vector along \mathbf{E} .
ϵ	dielectric permittivity [F m ⁻¹].
F	scattering potential [m ⁻²].
$f(\mathbf{s}, \mathbf{s}_0)$	scattering amplitude.
$\mathbf{F}(\mathbf{r}, t)$	flux density [W m ⁻²].
g	scattering anisotropy factor.
γ_2	excess kurtosis.
i	imaginary unit.
$I(\mathbf{r}, t, \mathbf{s})$	Specific intensity [W m ⁻² sr ⁻¹].
\mathbf{J}	current density [A m ⁻²].
\mathbf{k}	wave vector $ \mathbf{k} = 2\pi/\lambda$ [m ⁻¹].
L	thickness [m].
l_a	absorption mean free path [m].
λ	wavelength [m].
L_{eff}	effective thickness $L_{\text{eff}} = L + 2z_e$ [m].
l_s	scattering mean free path [m].
l'_s	reduced scattering mean free path [m].
μ	magnetic permeability [N A ⁻²].
μ_a	absorption coefficient [m ⁻¹].
μ_s	scattering coefficient [m ⁻¹].
μ'_s	reduced scattering coefficient [m ⁻¹].
n	refractive index $n = c/v$.
n_e	extraordinary refractive index in a birefringent crystal.

n_o	ordinary refractive index in a birefringent crystal.
OT	(reduced) optical thickness L/l'_s .
$p(\mathbf{s}, \mathbf{s}_0)$	phase function or scattering diagram.
ρ	radial coordinate in a cylindrical reference frame [m].
$R(\rho, t)$	time resolved reflectance [W m^{-2}].
$R(t)$	total time resolved reflectance [W].
\mathbf{S}	energy flux density [W m^{-2}].
\mathbf{s}	unit vector along \mathbf{S} .
σ_a	absorption cross section [m^2].
σ_s	scattering cross section [m^2].
σ_{tot}	extinction cross section [m^2].
\mathbf{r}	position vector [m].
t	time [s].
τ	decay constant of integrated transmittance/reflectance [s].
$T(\rho, t)$	time resolved transmittance [W m^{-2}].
$T(t)$	total time resolved transmittance [W].
$U(\mathbf{r}, t)$	average intensity [W m^{-2}].
$u(\mathbf{r}, t)$	energy density [J m^{-3}].
V	volume [m^3].
v	light speed in a medium $v = 1/\sqrt{\epsilon\mu}$ [m s^{-1}].
W	density of electromagnetic energy [J m^{-3}].
ω	angular frequency [rad s^{-1}].
$w^2(t)$	mean square width of a spatial profile [m^2].
z_e	extrapolation length [m].
z_{src}	isotropic point source depth [m].

” *There is something fascinating about science. One gets such wholesale returns of conjecture out of such a trifling investment of fact.*

— **Mark Twain**
(writer)

The presence of any transport dynamics in a complex, unknown medium represents an incredible opportunity to study its chemical and physical composition, as well as its microscopic structure or the nature of the interactions driving the transport process. Just by looking at the simple diffusive spreading of a droplet of ink in a glass of water we can unveil the temperature of the medium [1], while tracking human mobility patterns can give us valuable insight on how to prevent the outbreak of epidemics [2]. Much more conveniently than ink and humans, light (and waves in general) represents one of the most advantageous probes to investigate the properties of a host environment, being non-destructive, suitable to non-invasive experiments and easy to control and measure with a variety of techniques.

Even so, it can be extremely challenging to extract and take advantage of the wealth of information carried by the measured signal. As an example, conventional absorption spectroscopy requires clear and transparent materials, since it is necessary to know the path length through the medium to extract quantitative information. However, if the medium is turbid, the path length is no longer represented by the shortest distance from the light source to the spectrometer through the medium, and attenuation of signal can happen both because of scattering and absorption. On the other hand, the presence of scattering itself allows us to infer other structural properties of the sample that would be inaccessible otherwise, provided that appropriate models are available.

Given the relevance of wave transport, it's no surprise that a huge effort has been devoted to modeling its behavior, developing new theories and exploiting its applications. As a matter of fact, the very same theoretical models can often be applied seamlessly to study light propagation in an array of very different situations, ranging from a thin sheet of paper [3] to the whole atmosphere and oceans [4, 5], and from the internal structure of our body and brain [6, 7] to that of the Earth [8, 9], not to mention other planets and outer space [10, 11]. Yet, modeling light propagation can also help answer fundamental questions on the physics of transport [12–15], or be just a great source of fun and inspiration in the quest for ever more photorealistic rendering softwares [16, 17].

From a very general perspective, a typical transport problem consists of a medium with certain unknown (possibly tensor) properties $\mathbf{p}(\mathbf{r}, t)$, that we investigate with a

(monochromatic) light $X_{\text{in}}(\mathbf{r}, t, \mathbf{s})$ coming either from a source inside or outside the sample, with X representing some suitable radiometric quantity depending on the spatial coordinate \mathbf{r} , the direction \mathbf{s} , and time. It is usually assumed that $\mathbf{p}(\mathbf{r}, t) = \mathbf{p}(\mathbf{r})$ is static or quasi-constant in time, i.e., that any evolution of the optical properties of the medium occurs on a longer time scale than the propagation of light. Following interaction with the sample, some light $X_{\text{out}}(\mathbf{r}', t', \mathbf{s}')$ will eventually exit the sample. In this context, two tasks arise naturally as the *forward* and *inverse* problem. In the former case, we try to find a transfer function

$$f[X_{\text{in}}(\mathbf{r}, t, \mathbf{s}); \mathbf{p}(\mathbf{r})] \rightarrow X_{\text{out}}(\mathbf{r}', t', \mathbf{s}')$$

that allows us to predict X_{out} assumed that we know the properties of the sample and of the illumination. Radiative transfer theory, diffusion theory and Monte Carlo simulations are all examples of forward models for light transport. In the inverse case, which is most interesting for applications, we will rather seek a way to deduce $\mathbf{p}(\mathbf{r})$ assuming to have measured $X_{\text{out}}(\mathbf{r}', t', \mathbf{s}')$ or some part of it. This means finding

$$f^{-1}[X_{\text{in}}(\mathbf{r}, t, \mathbf{s}); X_{\text{out}}(\mathbf{r}', t', \mathbf{s}')] \rightarrow \mathbf{p}(\mathbf{r}).$$

These problems comprise the fundamental questions and motivation behind this work. The thesis is outlined as follows: in Chapter 2, we start by reviewing the main forward radiative transfer models and discuss their validity and range of application. The experimental aspects of ultrafast time-of-flight measurements are presented in Chapter 3, along with the development of a novel experimental configuration that enables the simultaneous investigation of both spatially and temporally-resolved transport. As illustrated by the ink droplet example, the inherently spatio-temporal concept of ‘spreading’ represents the most straightforward picture of the idea itself of propagation. Still, despite its simple representation, tracking light at the typical time and length scales associated to its transport dynamics poses several experimental challenges, requiring accurate calibration and validation of the measurement technique. Still, we argue that, almost by definition, transport cannot be really studied in its entirety disregarding either of these domains, nor by studying both of them, but separately. To support this claim, in Chapter 4 we bring evidence that current state-of-the-art, single-domain techniques are subject to pitfalls and shortcoming that prevent a correct optical characterization, or even the identification of novel transport regimes emerging in more extreme configurations. Most interestingly, in Chapter 5 we demonstrate how it is actually possible to take advantage of the subtle deviations from diffusion theory that we unveiled to implement a flexible and efficient inverse model based on a lookup-table routine and the gold-standard Monte Carlo method. In the last chapter, the origin of these deviations is elucidated by performing an extensive statistical characterization, which revealed the emergence of a well defined multiple scattering regime characterized by an effective diffusion constant that differs from that intrinsic to the material. This result, which could not be identified straightforwardly without the development of our new time-resolved imaging technique, is demonstrated under extremely general assumptions for the simple case of a homogeneous scattering slab, which represents the basic model in a number

of applications. As such, it could have far-reaching implications as it challenges our present interpretation of the link between the macroscopic and microscopic transport parameters of scattering media, and their optical characterization. In the last Section, further preliminary results are presented, revealing the presence of an even richer array of phenomena occurring in this simple system, ranging from anisotropic to anomalous diffusion and weakly self-similar transport.

” *From a long view of the history of mankind, seen from, say, ten thousand years from now, there can be little doubt that the most significant event of the 19th century will be judged as Maxwell’s discovery of the laws of electrodynamics. The American Civil War will pale into provincial insignificance in comparison with this important scientific event of the same decade.*

— **Richard P. Feynman**
(physicist)

The description of light transport in turbid media represents a challenging problem which has been studied for over a century, with applications in many different fields ranging from heat transfer [18] to astrophysics [19] and atmospheric sciences [20], to name a few. At the core of each of these fields, lies a condition of energy conservation expressed by the radiative transfer equation (RTE). The RTE is a balance equation for the quantity of energy

$$I(\mathbf{r}, t, \mathbf{s}) dt dS \cos \theta d\Omega$$

radiated per unit of time dt into a solid angle $d\Omega$ pointed by a unit vector \mathbf{s} forming an angle θ with the normal of a small area dS . The quantity $I(\mathbf{r}, t, \mathbf{s})$ [$\text{W m}^{-2} \text{sr}^{-1}$] is called the *specific intensity* and represents the main concept underlying the description of light transport in disordered media.

Despite the fact that the validity of the RTE has been validated across the most different fields (even beyond heat and light transport), its original derivation is purely phenomenological and was established simply as the expected energy balance when accounting for absorption and scattering. For most of the last 100 years, the RTE has thus been unrelated to the fundamental principles of classical electrodynamics. One of the main gaps was represented by the unclear relation between the specific intensity (a heuristically defined quantity) with the Poynting vector and Maxwell’s equations. The concept of an angular-dependent flow of energy expressed by the specific intensity has no direct counterpart in electromagnetic theory, where the Poynting vector field assumes a single value (and therefore a unique direction) in each point in space. The connection was finally established very recently both for the vector and scalar RTE [21, 22], considering a volume-averaged definition of the continuous Poynting vector field [23]. In the following section we will outline the main steps involved in the derivation of the scalar RTE from first principles. This serves two important purposes. Firstly, it allows us

to introduce all the optical parameters relevant to the description of radiative transport in disordered media. Secondly, it provides a complete review of all the approximations that are needed to obtain this fundamental equation, therefore providing clear insights on its validity domain.

2.1 From Maxwell's equations to Radiative transfer

2.1.1 Energy conservation

A relation of energy conservation arises naturally and directly from Maxwell's equations [24]. If we consider an isotropic medium with permeability μ and dielectric permittivity ϵ , Maxwell's equations can be expressed as

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \quad (2.1a)$$

$$\nabla \times \mathbf{B} = \epsilon\mu \frac{\partial \mathbf{E}}{\partial t} + \mu \mathbf{J} \quad (2.1b)$$

where \mathbf{E} , \mathbf{B} are the electric and the magnetic induction fields and \mathbf{J} is the density current. By multiplying these two equation by \mathbf{B} and \mathbf{E} respectively, and taking their difference we can rewrite

$$\frac{1}{\mu} \nabla \cdot (\mathbf{E} \times \mathbf{B}) = \frac{1}{\mu} \left(\epsilon\mu \mathbf{E} \cdot \frac{\partial \mathbf{E}}{\partial t} + \mathbf{B} \cdot \frac{\partial \mathbf{B}}{\partial t} \right) - \mathbf{E} \cdot \mathbf{J} \quad (2.2)$$

where we have used the vector identity $\nabla \cdot (\mathbf{a} \times \mathbf{b}) = \mathbf{b} \cdot (\nabla \times \mathbf{a}) - \mathbf{a} \cdot (\nabla \times \mathbf{b})$. Looking at equation (2.2), we can recognize the partial time derivative of the total electromagnetic energy

$$W(\mathbf{r}, t) = \frac{1}{2\mu} \left(\frac{1}{c^2} \mathbf{E} \cdot \mathbf{E} + \mathbf{B} \cdot \mathbf{B} \right) \quad [\text{J m}^{-3}] \quad (2.3)$$

and the divergence of the energy flux density

$$\mathbf{S} = \frac{1}{\mu} \mathbf{E} \times \mathbf{B}, \quad [\text{W m}^{-2}] \quad (2.4)$$

which allows us to interpret equation (2.2) as a continuity relation bounding the electromagnetic energy and its flux. The extra $\mathbf{E} \cdot \mathbf{J}$ term represents Joule's heating, expressing the rate of energy transfer from the field to the charges, i.e. dissipation of energy due to absorption. Energy dissipation in a isotropic medium is defined by Ohm's law as $\mathbf{J} = \omega\epsilon''\mathbf{E}$, where ω is the frequency of the electromagnetic wave and ϵ'' is the imaginary part of the permittivity $\epsilon = \epsilon'(\mathbf{r}) + i\epsilon''(\mathbf{r})$. Therefore, we can rewrite it in terms of the absorbed energy per unit volume

$$\frac{dP_{\text{abs}}}{dV} = \mathbf{E} \cdot \mathbf{J} = \omega\epsilon''\mathbf{E}^2, \quad [\text{W m}^{-3}] \quad (2.5)$$

which gives us the usual expression for Poynting's theorem

$$\frac{\partial W}{\partial t} + \frac{dP_{\text{abs}}}{dV} + \nabla \cdot \mathbf{S} = 0. \quad (2.6)$$

It should be emphasized that this relation holds at any point in space \mathbf{r} , provided that \mathbf{E} and \mathbf{B} are mutually orthogonal.

Equation (2.6) expresses energy conservation for time-harmonic electromagnetic fields. In a typical configuration, however, the period of an electromagnetic wave in the optical frequency range is several orders of magnitude larger than any experimental measurement time, and we are rather interested in *time-averaged* quantities. If we consider a plane wave solution to Maxwell's equations

$$\mathbf{E}(\mathbf{r}, t) = E_0 \mathbf{e} \exp(i\mathbf{k} \cdot \mathbf{r}) \exp(-i\omega t) \quad (2.7a)$$

$$\mathbf{B}(\mathbf{r}, t) = \sqrt{\epsilon\mu} E_0 \mathbf{b} \exp(i\mathbf{k} \cdot \mathbf{r}) \exp(-i\omega t) \quad (2.7b)$$

the time-averaged Poynting vector $\langle \mathbf{S} \rangle$ is simply given by

$$\langle \mathbf{S} \rangle \simeq \frac{1}{T} \int_0^T \frac{1}{\mu} [\mathbf{E}(\mathbf{r}, t) \times \mathbf{B}(\mathbf{r}, t)] dt = \frac{E_0^2}{2} \sqrt{\frac{\epsilon}{\mu}} \mathbf{s} \quad [\text{W m}^{-2}] \quad (2.8)$$

where $\mathbf{s} = \mathbf{e} \times \mathbf{b}$ are mutually orthogonal unitary vectors. This means that the validity of the time-averaged expression for the Poynting vector that we derived is limited to the *far-field*, where the electromagnetic fields propagate as a plane wave directed towards \mathbf{s} . Keeping in mind this assumption, we can analogously derive time-averaged expressions for the energy density and the absorbed power

$$\langle W \rangle = \frac{\epsilon}{2} E_0^2 = \sqrt{\epsilon\mu} \langle \mathbf{S} \rangle \cdot \mathbf{s} \quad [\text{J m}^{-3}] \quad (2.9)$$

$$\left\langle \frac{dP_{\text{abs}}}{dV} \right\rangle = \frac{1}{2} \omega \epsilon'' E_0^2 = \sqrt{\frac{\mu}{\epsilon}} \omega \epsilon'' \langle \mathbf{S} \rangle \cdot \mathbf{s} \quad [\text{W m}^{-3}] \quad (2.10)$$

where $\frac{1}{\sqrt{\epsilon\mu}} = v$ is the speed of light in the medium. We can now use the time-averaged expressions (2.8), (2.9) and (2.10) to rewrite

$$\frac{1}{v} \frac{\partial \langle \mathbf{S}(\mathbf{r}) \rangle \cdot \mathbf{s}}{\partial t} + \left\langle \frac{dP_{\text{abs}}}{dV} \right\rangle + \nabla \cdot \langle \mathbf{S}(\mathbf{r}) \rangle = 0, \quad (2.11)$$

which is the time-averaged expression of equation (2.6) and represents the conservation of energy flux along the direction of the Poynting vector \mathbf{s} . Of course, energy conservation must hold along any arbitrary direction \mathbf{s}_j , and we can rewrite

$$\frac{1}{v} \frac{\partial \langle \mathbf{S}(\mathbf{r}) \rangle \cdot \mathbf{s}_j}{\partial t} + \left\langle \frac{dP_{\text{abs}}}{dV} \right\rangle (\mathbf{s} \cdot \mathbf{s}_j) + \mathbf{s}_j \cdot \nabla (\langle \mathbf{S}(\mathbf{r}) \rangle \cdot \mathbf{s}_j) = 0, \quad (2.12)$$

where we have used the relation

$$\nabla \cdot (\langle \mathbf{S}(\mathbf{r}) \rangle \cdot \mathbf{s}_j) \mathbf{s}_j = \mathbf{s}_j \cdot \nabla (\langle \mathbf{S}(\mathbf{r}) \rangle \cdot \mathbf{s}_j). \quad (2.13)$$

Equation (2.12) expresses the fact that energy conservation is rotationally invariant, i.e. that power is conserved irrespective of the angle that a detector holds with the power flux. The total power measured experimentally by a detector of area A placed at \mathbf{r} with surface normal \mathbf{n} can be therefore expressed as $P(\mathbf{r}) = \int_A \langle \mathbf{S}(\mathbf{r}') \rangle \cdot \mathbf{n} dS'$. Finally, in the case of a non-absorbing medium ($dP_{\text{abs}}/dV = 0$) under continuous illumination ($\langle \partial W / \partial t \rangle = 0$) which contains no sources, the conservation of energy simply states that $\nabla \cdot \langle \mathbf{S} \rangle = 0$ or, alternatively, that the averaged total flux of energy $\int_{\Sigma} \langle \mathbf{S} \rangle \cdot \mathbf{n} dS$ through any closed surface Σ is zero.

2.1.2 Optical parameters of a particle

In order to describe light propagation in a turbid medium we must take into account the presence of inhomogeneities. Let us therefore introduce a spatially varying index of refraction $n(\mathbf{r})$ representing an isolated scattering particle of volume V and arbitrary shape embedded in a host material of refractive index n_0 . From now on, we will always consider the host and scattering media to be non-magnetic. In this case, in absence of charges and currents (i.e. $\mathbf{J} = 0$), we can take the time derivative of equation (2.1b) and substitute the expression for $\partial B / \partial t$ obtaining

$$-\nabla \times (\nabla \times \mathbf{E}) - \frac{n^2(\mathbf{r})}{c^2} \frac{\partial^2 \mathbf{E}}{\partial t^2} = -\frac{n^2(\mathbf{r})}{c^2} \frac{\partial^2 \mathbf{E}}{\partial t^2} + \nabla^2 \mathbf{E} - \nabla(\nabla \cdot \mathbf{E}) = 0 \quad (2.14)$$

where we have used the identity $\nabla \times (\nabla \times \mathbf{a}) = \nabla(\nabla \cdot \mathbf{a}) - \nabla^2 \mathbf{a}$. Assuming a time-harmonic dependence of the field (2.7a), equation (2.14) becomes

$$\nabla^2 \mathbf{E}(\mathbf{r}) + \frac{n_0^2 \omega^2}{c^2} \mathbf{E} = \frac{n_0^2 \omega^2}{c^2} \left(\frac{n^2(\mathbf{r})}{n_0^2} - 1 \right) \mathbf{E}(\mathbf{r}) + \nabla(\nabla \cdot \mathbf{E}(\mathbf{r})) \quad (2.15)$$

where $n_0 \omega / c = 2\pi / \lambda = k$ represents the wavenumber of the propagating wave of angular frequency ω and wavelength λ . The term $F(\mathbf{r}) = k^2 (n^2(\mathbf{r}) / n_0^2 - 1)$ is usually referred to as the *scattering potential*, and vanishes for \mathbf{r} outside V . Equation (2.15) represents the full scattering problem for the electric field vector, including its change in polarization due to the source term $\nabla(\nabla \cdot \mathbf{E}(\mathbf{r}))$, which couples the cartesian components of \mathbf{E} . If we assume that $n(\mathbf{r})$ varies slowly on length scales comparable to λ , or we decide to ignore polarization effects altogether, we can neglect the coupling term and obtain the (uncoupled) *scalar* differential equations

$$\nabla^2 E(\mathbf{r}) + k^2 E(\mathbf{r}) = F(\mathbf{r}) E(\mathbf{r}) \quad (2.16)$$

which can be more easily solved for each component.

The solution to equation (2.16) for any point outside the scatterer can be written as the combination of an incident and a scattered field $E(\mathbf{r}) = E_{\text{inc}}(\mathbf{r}) + E_{\text{sc}}(\mathbf{r})$ (Figure 2.1), where E_{inc} corresponds to the value of the field in the absence of the particle, while

$$E_{\text{sc}}(\mathbf{r}) = \int_V F(\mathbf{r}') E(\mathbf{r}') G(\mathbf{r}, \mathbf{r}') d^3 \mathbf{r}', \quad (2.17)$$




Fig. 2.1.: T-Matrix calculation of the electric field on the xz -plane following single scattering from a 3D dielectric sphere ($d = 250$ nm, $n = 2.7$, $n_0 = 1.5$). The incoming field is a monochromatic plane wave ($\lambda = 532$ nm) propagating from left to right.

with $G(\mathbf{r}, \mathbf{r}') = G(|\mathbf{r} - \mathbf{r}'|) = \exp(ik|\mathbf{r} - \mathbf{r}'|)/4\pi|\mathbf{r} - \mathbf{r}'|$ being the free-space outgoing Green function. If we assume that our detector is placed in the far field of the particle ($r \gg r'$), we can approximate $|\mathbf{r} - \mathbf{r}'| \sim r - \mathbf{s} \cdot \mathbf{r}'$ with \mathbf{s} being the unit vector along \mathbf{r} , in which case the Green function factorizes into $G(|\mathbf{r} - \mathbf{r}'|) \sim \exp(ikr) \exp(-iks \cdot \mathbf{r}')/4\pi r$. Combining this equation with the plane wave incident field $E_{\text{inc}}(\mathbf{r}) = E_0 \exp(iks \cdot \mathbf{r})$ propagating with $\mathbf{k}_{\text{inc}} = k\mathbf{s}_0$ along the direction \mathbf{s}_0 , we obtain the expression for the scattered electric field as

$$E_{\text{sc}}(\mathbf{r}) = E_0 f(\mathbf{s}, \mathbf{s}_0) \frac{e^{ikr}}{r} \quad (2.18)$$

where we have introduced the *scattering amplitude* $f(\mathbf{s}, \mathbf{s}_0)$ as

$$f(\mathbf{s}, \mathbf{s}_0) = \frac{1}{4\pi} \int_V F(\mathbf{r}') \frac{E(\mathbf{r}')}{|E_0|} e^{-iks \cdot \mathbf{r}'} d^3 r' \quad (2.19)$$

which is defined with respect to the incident direction \mathbf{s}_0 and independent on the amplitude of E_{inc} . Finally, this allows us to express the components of time-averaged Poynting vector $\langle \mathbf{S}_{\text{sc}} \rangle$ associated with the incident and scattered field as

$$\langle \mathbf{S}_{\text{inc}} \rangle = \frac{E_0^2 \epsilon v}{2} \mathbf{s}_0 \quad (2.20)$$

$$\langle \mathbf{S}_{\text{sc}} \rangle = \frac{E_0^2 \epsilon v}{2} \frac{|f(\mathbf{s}, \mathbf{s}_0)|^2}{r^2} \mathbf{s} = |\langle \mathbf{S}_{\text{inc}} \rangle| \frac{|f(\mathbf{s}, \mathbf{s}_0)|^2}{r^2} \mathbf{s}. \quad (2.21)$$

Note that the full time-averaged Poynting vector $\langle \mathbf{S} \rangle$ will include an additional term due to the interference between scattered and incident fields.

Using equation (2.18) we are now able to define several common quantities which refer directly to the properties of the particle and that eventually determine the way light propagates through disordered, opaque media.

Let us consider the amount of energy lost by the interaction of the incident light on the particle due to absorption. We can write an expression for \bar{P}_{abs} as

$$\bar{P}_{\text{abs}} = \int_V \left\langle \frac{dP_{\text{abs}}}{dV} \right\rangle dV = \frac{\omega}{2} \int_V \epsilon'(\mathbf{r}) |E(\mathbf{r})|^2 dV \quad [\text{W}] \quad (2.22)$$

representing the amount of energy lost per second due to absorption. Normalizing this power by the rate at which energy impinges on the particle, we obtain an *absorption cross-section*

$$\sigma_a = \frac{\bar{P}_{\text{abs}}}{|\langle \mathbf{S}_{\text{inc}} \rangle|} = \frac{k}{2E_0^2 \epsilon} \int_V \epsilon'(\mathbf{r}) |E(\mathbf{r})|^2 dV \quad [\text{m}^2] \quad (2.23)$$

which depends solely on the material properties of the particle and its geometry. A *scattering cross-section* is similarly obtained as

$$\sigma_s = \frac{\bar{P}_{\text{sc}}}{|\langle \mathbf{S}_{\text{inc}} \rangle|} = \int_V \frac{\nabla \cdot \langle \mathbf{S}_{\text{sc}} \rangle}{|\langle \mathbf{S}_{\text{inc}} \rangle|} dV = \int_S \frac{\langle \mathbf{S}_{\text{sc}} \rangle \cdot \mathbf{n}}{|\langle \mathbf{S}_{\text{inc}} \rangle|} dS. \quad [\text{m}^2] \quad (2.24)$$

Introducing the far-field expressions for $\langle \mathbf{S}_{\text{sc}} \rangle$ and $\langle \mathbf{S}_{\text{inc}} \rangle$ we obtain

$$\sigma_s = \int_S |f(\mathbf{s}, \mathbf{s}_0)|^2 \frac{dS}{r^2} = \int_{4\pi} |f(\mathbf{s}, \mathbf{s}_0)|^2 d\Omega, \quad [\text{m}^2] \quad (2.25)$$

where we have substituted dS/r^2 with the solid angle $d\Omega$ assuming that the center of integration is positioned at the center of the particle.

The sum of the absorption and scattering cross-sections is usually referred to as the total or extinction cross-section

$$\sigma_{\text{tot}} = \sigma_a + \sigma_s, \quad [\text{m}^2] \quad (2.26)$$

which is a particularly relevant microscopic quantity since it depends only on the incident direction of propagation and can be measured experimentally in the far-field, providing important information on the microscopic properties of a scattering particle through the *optical theorem* $\sigma_{\text{tot}} = 4\pi \text{Im}\{f(\mathbf{s}, \mathbf{s}_0)\}/k$.

It is customary to introduce a separate function to refer directly to the square modulus of the scattering amplitude. This function, misleadingly referred to as the *phase function* (it bears no relation with the phase of the electromagnetic wave), is conveniently normalized by the total cross-section

$$p(\mathbf{s}, \mathbf{s}_0) = \frac{1}{\sigma_{\text{tot}}} |f(\mathbf{s}, \mathbf{s}_0)|^2 \quad (2.27)$$

and therefore can be interpreted statistically as the probability distribution for light incident on the particle from direction \mathbf{s}_0 to be scattered in direction \mathbf{s} . Using the definition of the scattering cross-section (2.24) we obtain the following relation

$$\int_{4\pi} p(\mathbf{s}, \mathbf{s}_0) d\Omega = \frac{\sigma_s}{\sigma_{\text{tot}}} = \varpi_0, \quad (2.28)$$




Fig. 2.2.: Comparison between Henyey-Greenstein and Mie phase function ($g = 0.6034$). Mie data taken from the Mie scattering calculator [26], assuming a 3D dielectric sphere ($d = 250$ nm, $n = 2.7$, $n_0 = 1.5$), and a monochromatic incoming plane wave ($\lambda = 532$ nm) propagating from left to right.

where ϖ_0 is usually called the *albedo* and becomes unity for a non-absorbing particle. The phase function is an extremely complex function which takes into account all the possible interference effects that occur inside the particle, and can only be solved analytically for very ideal and simple shapes such as a sphere (e.g., using Mie theory [25], see Fig. 2.2). However, when considering large assemblies of statistically equivalent scatterers, it is customary to use approximated forms. If the scatterers are randomly oriented, we can assume that the scattering phase function is independent on the direction of propagation, i.e. $p(\mathbf{s}, \mathbf{s}_0) = p(\mathbf{s} \cdot \mathbf{s}_0)$. In this case, a general phase function can be defined as an expansion over Legendre polynomials

$$p(\mathbf{s} \cdot \mathbf{s}_0) = \sum_{l=0}^{\infty} a_l P_l(\mathbf{s} \cdot \mathbf{s}_0). \quad (2.29)$$

A simple choice is that of taking $a_l = (2l - 1)g^l$ which leads to the widely used Henyey-Greenstein phase function [27]

$$p_{\text{HG}}(\mathbf{s} \cdot \mathbf{s}_0) = \frac{\varpi_0}{4\pi} \frac{1 - g^2}{(1 + g^2 - 2g\mathbf{s} \cdot \mathbf{s}_0)^{3/2}} \quad (2.30)$$

which depends solely on the *scattering anisotropy factor* $g \in [-1, 1]$

$$g = \langle \mathbf{s} \cdot \mathbf{s}_0 \rangle = \langle \cos \theta \rangle = \frac{\int_{4\pi} p(\mathbf{s}, \mathbf{s}_0) \mathbf{s} \cdot \mathbf{s}_0 d\Omega}{\int_{4\pi} p(\mathbf{s}, \mathbf{s}_0) d\Omega}, \quad (2.31)$$

a general parameter expressing the scattering directionality from completely forward ($g = 1$) to completely backwards ($g = -1$), with $g = 0$ representing the isotropic scattering case.

2.1.3 Multiple scattering

Let us now consider a collection of N particles with equal optical properties. In the far field, we can write the scattered field at \mathbf{r} as the summation

$$E_{sc}(\mathbf{r}) = \sum_{i=1}^N E_{i,sc}(\mathbf{r}) = |E_0| \sum_{i=1}^N f_i(\mathbf{s}_i, \mathbf{s}_0) \frac{e^{ik|\mathbf{r}-\mathbf{r}_i|}}{|\mathbf{r}-\mathbf{r}_i|}, \quad (2.32)$$

where each $f_i(\mathbf{s}_i, \mathbf{s}_0)$ depends on the *total* field $E(\mathbf{r}_i) = E_{inc}(\mathbf{r}_i) + E_{sc}(\mathbf{r}_i)$ at \mathbf{r}_i as expressed in equation (2.19). The intensity $|E|^2$, which is needed to express the energy density flow $\langle \mathbf{S} \rangle$, will contain a contribution from each particle as well

$$|E(\mathbf{r})|^2 = \left(E_{inc}(\mathbf{r}) + \sum_{i=1}^N E_{i,sc}(\mathbf{r}) \right) \left(E_{inc}(\mathbf{r}) + \sum_{i=1}^N E_{i,sc}(\mathbf{r}) \right)^* \quad (2.33)$$

and in all practical cases cannot be solved analytically even when ignoring depolarization effects. Introducing equation (2.33) into the expression for $\langle \mathbf{S} \rangle$ we can rewrite

$$\langle \mathbf{S}(\mathbf{r}) \rangle = \langle \mathbf{S}_{inc}(\mathbf{r}) \rangle + \sum_{i=1}^N \langle \mathbf{S}_{sc}(\mathbf{r}) \rangle_i + \sum_{\substack{i,j=1 \\ i \neq j}}^N \langle \mathbf{S}_{sc}(\mathbf{r}) \rangle_{ij} + \dots, \quad (2.34)$$

where $\langle \mathbf{S}_{sc} \rangle_i$ is the contribution due to particle i , $\langle \mathbf{S}_{sc} \rangle_{ij}$ results from the interference from particle i and j , and all higher order terms have not been written explicitly, as the series contains as many terms as there are combinations between N particles. In order to simplify the problem we will assume that the wavelength of light is much smaller than the typical distances involved in the problem, and neglect all interference effects due to the scattering between particles. Each particle will then contribute independently to the energy density flux

$$\langle \mathbf{S}(\mathbf{r}) \rangle = \langle \mathbf{S}_{inc}(\mathbf{r}) \rangle + \sum_{i=1}^N \langle \mathbf{S}_{sc}(\mathbf{r}) \rangle_i, \quad (2.35)$$

where the scattering components can be rewritten as

$$\langle \mathbf{S}_{sc}(\mathbf{r}) \rangle_i = \sigma_{tot} \langle \mathbf{S}_{inc}(\mathbf{r}_i) \rangle \cdot \mathbf{s} \frac{p(\mathbf{s}_i, \mathbf{s})}{|\mathbf{r}-\mathbf{r}_i|^2} \mathbf{s}_i \quad (2.36)$$

using the definitions of $\langle \mathbf{S}_{inc} \rangle$ and of the phase function.

In many practical situations we do not know the number of particles N nor their sizes or shapes. However, if we assume that they all have the same average radius R and random orientation we can introduce few average quantities with statistical significance such as the *scattering coefficient* and the *absorption coefficient* (or their associated mean free paths)

$$\mu_s = n\sigma_s = 1/l_s \quad [\text{m}^{-1}] \quad (2.37a)$$

$$\mu_a = n\sigma_a = 1/l_a \quad [\text{m}^{-1}] \quad (2.37b)$$

where n is the number density of particles in the element of volume. Finally, to take into account the scattering anisotropy factor g , another quantity is introduced to relate the scattering coefficient of an anisotropic particle to that of an isotropic one

$$\mu'_s = \mu_s(1 - g) = 1/l'_s \quad [\text{m}^{-1}] \quad (2.38)$$

which is called the *reduced scattering coefficient*. Its reciprocal, the reduced scattering mean free path or transport mean free path l'_s , represents the distance that light needs to travel before it loses every residual correlation with its original direction of propagation. This distance diverges asymptotically as $g \rightarrow 1$, while it is equivalent to l_s for $g = 0$, consistently with the definition of isotropic scattering. As we will see in the following, this relation expresses the fact that, under the diffusive approximation, the transport of light in a medium characterized by scattering parameters μ_s and $g \neq 0$ can be mapped identically onto that of a system with μ'_s and $g = 0$. This degeneracy is usually referred to as the *similarity relation* [25, 28].

It is important to stress the statistical nature of the parameters that we introduced to describe radiative transport. Both the scattering/absorption coefficients and our choice of a phase function are basically independent of the actual size and properties of each individual particle, and make sense only from a statistical point of view. In most applications, though, this statistical approach and the loss of wave properties that characterizes multiple scattering in the radiative transfer framework does not represent a problem. The reason is that often, both in fundamental research and applications, statistical properties represent actually the most meaningful information, and it would be pointless to fully resolve the extremely complex (but deterministic) way in which light propagates in a very specific configuration of scatterers rather than in any other statistically equivalent one. As a matter of fact, in many practical cases we may even want to have our scattering sample moving or, equivalently, average its transport properties over many different regions in order to converge more efficiently towards its average properties, especially for heterogeneous media [29]. Even where exact techniques are available to solve numerically the electromagnetic problem, such as the T-Matrix method, random orientation-averaging is frequently applied to obtain meaningful average properties [30]. A useful analogy is that of a pile of sand: being able to retrieve statistical properties such as their local packing density, distribution of grain sizes or positional and angular correlations functions, is much more informative and desirable than obtaining the exact number of grains, let alone their individual shapes, positions and orientation.

That being said, the fundamental nature of multiple wave scattering is indeed deterministic, which must be taken into account to explain a number of different phenomena and applications which are worth mentioning, even if briefly. Among the most striking examples, phase-conjugation and its time-reversal applications have been investigated since the early 80's [31, 32]. Moreover, there are actually cases where we might be interested in the existence of specific, non-general solutions to the problem of transport that are applicable only to a single specimen. Notable examples include the




Fig. 2.3.: From left to right: representation of the Poynting vector field $\langle \mathbf{S}(\mathbf{r}') \rangle$ evaluated at $\mathbf{r}' \in \delta V$. Averaging over the differential volume δV we obtain a volume-averaged flow of energy $\|\mathbf{S}(\mathbf{r})\|_v$ pointing at an average direction $\langle \mathbf{s}_r \rangle_v$, with different projections along arbitrary directions \mathbf{s}_j .

problem of perfect focusing and the control of transmission through scattering media [33, 34], as well as the enhancement of photonic devices through the introduction of intentional disorder [35] or the realization of physically unclonable keys for secure authentication [36]. On a more fundamental side, well-known phenomena such as coherent backscattering [37] and Anderson localization [38] take place in scattering, disordered media and yet are entirely based on interference effects. The same holds for the concept of correlated disorder, which can give rise to a wide array of intriguing interference effects ranging from non-iridescent structural coloration, to bang-gap formation and transparency [15, 39–41]. Finally, the passive retrieval of the time-dependent Green function has been also recently demonstrated in strongly scattering media, based on the measurement of the mutual coherence of an incoherent excitation at two different points and times [42]. All the above examples use the fact that light transport in scattering materials is deterministic, as opposed to the stochastic nature of the RTE.

2.1.4 Specific intensity, average intensity and flux

In order to derive the RTE, we consider the time-averaged Poynting vector $\langle \mathbf{S} \rangle = S\mathbf{s}$ and take its average over a small volume δV [22, 23]. Most of the approximations involved on the derivation of the RTE are based on the ‘smallness’ of δV , and its role will be further discussed at the end of the section.

For each point $\mathbf{r}' \in \delta V$ the Poynting vector will have a well defined value $\langle \mathbf{S}(\mathbf{r}') \rangle$ pointing along a certain direction \mathbf{s} (Fig. 2.3). The resulting integral can have contributions along any direction \mathbf{s}_j , each with a normalized weight

$$w_r(\mathbf{s}_j) = \frac{1}{\delta V \|\mathbf{S}(\mathbf{r})\|_v} \int_{\delta V} S(\mathbf{r} - \mathbf{r}') (\mathbf{s}' \cdot \mathbf{s}_j) d^3 r', \quad [\text{sr}^{-1}] \quad (2.39)$$

where $\|\mathbf{S}(\mathbf{r})\|_v$ is the magnitude of volume-averaged flow

$$\|\mathbf{S}(\mathbf{r})\|_v = \frac{1}{\delta V} \int_{\delta V} S(\mathbf{r} - \mathbf{r}') d^3 r'. \quad (2.40)$$

We can now express the average direction of energy flow

$$\langle \mathbf{s}_r \rangle_v = \frac{1}{4\pi} \int_{4\pi} w_r(\mathbf{s}_j) \mathbf{s}_j d\Omega \quad (2.41)$$

and finally write the volume-averaged energy flow as

$$\langle \mathbf{S}(\mathbf{r}) \rangle_v = \|\mathbf{S}(\mathbf{r})\|_v \langle \mathbf{s}_r \rangle_v, \quad (2.42)$$

which is the quantity that will be used to establish a connection with the concept of specific intensity and other radiometric quantities.

As we mentioned at the beginning of the chapter, the specific intensity represents the amount of power per unit area that flows in a certain direction defined by a unit solid angle. In terms of the average energy flow defined in equation (2.42), we can define it as

$$I(\mathbf{r}, t, \mathbf{s}) = \frac{1}{4\pi} \|\mathbf{S}(\mathbf{r})\|_v w_r(\mathbf{s}) \quad [\text{W m}^{-2} \text{ sr}^{-1}] \quad (2.43)$$

where $w_r(\mathbf{s})$ expresses the probability of the averaged flow to point in direction \mathbf{s} as shown in equation (2.39). We must note that the specific intensity, as well as all other radiometric quantities, are inherently referred to a certain frequency interval: we will in general omit such dependence and consider quasi-monochromatic light.

Starting from the specific intensity, another relevant quantity that we can build is the average intensity at a point \mathbf{r} , defined as the volume average of the Poynting vector integrated over all directions

$$U(\mathbf{r}, t) = \int_{4\pi} I(\mathbf{r}, t, \mathbf{s}) d\Omega = \frac{1}{4\pi} \|\mathbf{S}(\mathbf{r})\|_v \int_{4\pi} w_r(\mathbf{s}) d\Omega = \|\mathbf{S}(\mathbf{r})\|_v \quad [\text{W m}^{-2}] \quad (2.44)$$

which is equal to the magnitude of the volume averaged energy flow. Connected to it, a quantity that is commonly used in the radiative transfer formalism is the energy density $u(\mathbf{r}, t)$. A convenient way to define it is to consider the definition of energy density (2.10) that we obtained directly from Maxwell's equations and relate it to the volume averaged Poynting vector

$$u(\mathbf{r}, t) = \frac{1}{v\delta V} \int_{\delta V} |\langle \mathbf{S}(\mathbf{r} - \mathbf{r}') \rangle| d^3r' = \frac{1}{v} \|\mathbf{S}(\mathbf{r})\|_v = \frac{U(\mathbf{r}, t)}{v}. \quad [\text{J m}^{-3}] \quad (2.45)$$

Neither $U(\mathbf{r}, t)$ nor $u(\mathbf{r}, t)$ contain any information regarding the average direction of propagation, though. The main quantity that is connected to the overall flux of energy is the flux density

$$\mathbf{F}(\mathbf{r}, t) = \int_{4\pi} I(\mathbf{r}, t, \mathbf{s}) \mathbf{s} d\Omega = \frac{1}{4\pi} \|\mathbf{S}(\mathbf{r})\|_v \int_{4\pi} w_r(\mathbf{s}) \mathbf{s} d\Omega = \|\mathbf{S}(\mathbf{r})\|_v \langle \mathbf{s}_r \rangle_v \quad [\text{W m}^{-2}] \quad (2.46)$$

which is a vector with magnitude and direction of the average flow of electromagnetic energy, which we have already encountered in equation (2.42).

2.1.5 The radiative transfer equation

The starting point to derive the radiative transfer equation (RTE) is represented by the equation of energy conservation expressed by Poynting's theorem (2.6). Even in the presence of a time dependence (i.e., if we have an intensity-modulated source), provided that it is much slower than the optical oscillation period, we can resort to our time-averaged expression (2.12) for energy conservation along any arbitrary direction \mathbf{s}_j . Let us consider the differential volume δV as the one depicted in Figure 2.3, containing N particles with absorption cross-section σ_a and scattering cross section σ_s . If we integrate equation (2.12) over δV we obtain the following expression

$$\begin{aligned} \frac{1}{v} \frac{\partial}{\partial t} \int_{\delta V} (\mathbf{s} \cdot \mathbf{s}_j) S(\mathbf{r} - \mathbf{r}') d^3 r' \\ + \int_{\delta V} (\mathbf{s} \cdot \mathbf{s}_j) \left\langle \frac{dP_{\text{abs}}(\mathbf{r} - \mathbf{r}')}{dV} \right\rangle d^3 r' \\ + \int_{\delta V} (\mathbf{s} \cdot \mathbf{s}_j) \mathbf{s}_j \cdot \nabla_{\mathbf{r}'} S(\mathbf{r} - \mathbf{r}') d^3 r' = 0 \quad (2.47) \end{aligned}$$

which, after several approximations, can be connected to the RTE in terms of the specific intensity. An overview of the steps needed to derive each one of these terms can be found in Appendix A.1, here we will limit ourselves to give their final respective representation as the

- Volume-averaged change in energy density

$$\frac{1}{v} \frac{\partial}{\partial t} \int_{\delta V} (\mathbf{s} \cdot \mathbf{s}_j) S(\mathbf{r} - \mathbf{r}') d^3 r' \rightarrow \frac{1}{v} \frac{\partial}{\partial t} \|S(\mathbf{r})\|_v w_r(\mathbf{s}_j) \delta V \quad (2.48)$$

- Volume-averaged absorbed power

$$\int_{\delta V} (\mathbf{s} \cdot \mathbf{s}_j) \left\langle \frac{dP_{\text{abs}}(\mathbf{r} - \mathbf{r}')}{dV} \right\rangle d^3 r' \rightarrow N \sigma_a \|S(\mathbf{r})\|_v w_r(\mathbf{s}_j) \quad (2.49)$$

- Volume-averaged change in energy flow

$$\begin{aligned} \int_{\delta V} (\mathbf{s} \cdot \mathbf{s}_j) \mathbf{s}_j \cdot \nabla_{\mathbf{r}'} S(\mathbf{r} - \mathbf{r}') d^3 r' \rightarrow \mathbf{s}_j \cdot \nabla \|S(\mathbf{r})\|_v w_r(\mathbf{s}_j) \delta V \\ + N \sigma_s \|S(\mathbf{r})\|_v w_r(\mathbf{s}_j) - N \sigma_{\text{tot}} \int_{4\pi} \|S(\mathbf{r})\|_v w_r(\mathbf{s}_j') p(\mathbf{s}_j, \mathbf{s}_j') d\Omega'. \quad (2.50) \end{aligned}$$

The dependence on the arbitrary volume δV can be removed introducing the density of particles $n = N/\delta V$, which we already used in the definitions for the scattering and absorption coefficients (2.37b). Finally, if we identify the specific intensity as

$\|\mathbf{S}(\mathbf{r})\|_v w_r(\mathbf{s})/4\pi$ (equation (2.43)) we can recast equation (2.47) into the usual equation for radiative transfer

$$\frac{1}{v} \frac{\partial I(\mathbf{r}, t, \mathbf{s})}{\partial t} = -\mathbf{s} \cdot \nabla I(\mathbf{r}, t, \mathbf{s}) - (\mu_s + \mu_a) I(\mathbf{r}, t, \mathbf{s}) + \mu_{\text{tot}} \int_{4\pi} I(\mathbf{r}, t, \mathbf{s}') p(\mathbf{s}, \mathbf{s}') d\Omega'. \quad (2.51)$$

Clearly, the RTE represents an energy balance for the flow of energy, stating that the specific intensity at \mathbf{r} at time t pointed towards direction \mathbf{s} will vary due to, respectively, the energy flowing through its boundaries $\mathbf{s} \cdot \nabla I(\mathbf{r}, t, \mathbf{s})$, the losses from absorption and scattering $(\mu_s + \mu_a) I(\mathbf{r}, t, \mathbf{s})$ and the gain due to scattering from any direction \mathbf{s}' into direction \mathbf{s} : $\mu_{\text{tot}} \int_{4\pi} I(\mathbf{r}, t, \mathbf{s}') p(\mathbf{s}, \mathbf{s}') d\Omega'$. The last term missing from this energy balance is the source term, which can assume different forms depending on the specific problem. If we consider to have a certain amount of energy delivered to the system, we can model it as an additional gain term $Q(\mathbf{r}, t, \mathbf{s})$ emitting a certain amount of energy in direction \mathbf{s} .

An important property of the RTE is that it is invariant under certain transformations. Two of these similarity relations are particularly relevant. If $I(\mathbf{r}, t, \mathbf{s})$ is a solution of equation (2.51) for a medium with $\mu_{\text{tot}} = \mu_s + \mu_a$ and phase function $p(\mathbf{s}, \mathbf{s}_0)$, then we can scale the specific intensity as

$$\tilde{I}(\tilde{\mathbf{r}}, \tilde{t}, \mathbf{s}) = \left(\frac{\tilde{\mu}_{\text{tot}}}{\mu_{\text{tot}}} \right)^3 I(\mathbf{r}, t, \mathbf{s}) \quad (2.52)$$

with $\tilde{\mathbf{r}} = \mathbf{r} \mu_{\text{tot}} / \tilde{\mu}_{\text{tot}}$ and $\tilde{t} = t \mu_{\text{tot}} / \tilde{\mu}_{\text{tot}}$. This conveniently allows to use the results obtained in a given geometry for any other scaled geometry, as long as the albedo $\varpi_0 = \mu_s / \mu_{\text{tot}}$ and $p(\mathbf{s}, \mathbf{s}_0)$ remain constant [43]. Perhaps even more important, another relation exist linking the solutions for a medium with and without absorption, namely

$$I(\mathbf{r}, t, \mathbf{s}) = I(\mathbf{r}, t, \mathbf{s} | \mu_a = 0) e^{-\mu_a v t}, \quad (2.53)$$

which is particularly relevant when attempting to solve the RTE numerically, since it allows to calculate only one solution per each (μ_s, g) pair, to which any μ_a dependence can be applied *a posteriori* [44].

It is worth summarizing the main approximations taken to derive the RTE from Maxwell's equations. The far-field approximation is perhaps the strongest approximation needed to derive formally the RTE and the definition of the specific intensity itself. It allows us to consider the contribution from each scatterer as an outgoing spherical wave, with the electric and magnetic field always mutually orthogonal. Connected to this, we have always dealt with additive intensities rather than fields, i.e. we have completely neglected any interference effect. Analogously, this has allowed us to sum incoherently the scattering and absorption cross-sections, assuming an *independent* scattering regime that is strictly valid only for very low concentrations of scatterers. Secondly, we assumed that our sampling volume δV was large enough to contain a statistically representative amount of scatterers, but yet so small to contribute negligibly to the scattered flow from the rest of the sample volume. This also allowed us to substitute the Poynting vector within δV with its volume-averaged value. Finally, in order to use the scalar wave

equation, we neglected the term $\nabla(\nabla \cdot \mathbf{E})$ that couples different polarization components in equation (2.15). We must note that this approximation is greatly alleviated when one averages over multiple realization of disorder. Moreover it is pertinent to the derivation of the scalar RTE only, since polarization can be accounted for in the full derivation of the vector RTE, and indeed it has been shown to play a non-trivial role in determining transport properties [45].

Despite all the approximations that are involved in a rigorous derivation of the RTE from electromagnetic theory, the radiative transfer framework that we described has proven to be extremely robust against its own defining assumptions. As a matter of fact, the extent of its validity range and flexibility goes well beyond our expectations, both in terms of the variety of fields and the physical configurations to which it applies. Few notable examples include the fact that the RTE is known to describe light transport accurately also for samples whose thickness is comparable with the wavelength [46], or with a dense packing of particles, where coherent (dependent) scattering needs to be taken into account and the far-field approximation clearly breaks down. When finite-size scatterers are packed together at high density, their spatial arrangement is not entirely random, and a short-range order emerges akin to that of molecules in a liquid. As it turns out, it is possible to map this dependent scattering problem back onto an independent regime with effective parameters that are modulated according to the spatial correlations of the sample, as described by the structure factor of the spatial arrangement [47, 48].

2.1.6 Solving the RTE: Monte Carlo method

Despite the multiple approximations that we applied, the RTE is a integro-differential equation that is exceedingly difficult to solve explicitly. Several approaches exist to approximate its solution (see [7, 49] and references therein for a comprehensive discussion), one of the main being the spherical harmonics expansion of the specific intensity $I(\mathbf{r}, t, \mathbf{s})$, often referred to as the P_N approximation by the number N of terms at which the expansion is truncated. Promising progress is recently being reported in this direction, exploiting a newly developed rotated reference frame method [50, 51], though mainly limited to a steady-state description of radiative transport.

A different approach is that represented by the Monte Carlo (MC) method, where the conservation of energy expressed by equation (2.51) is enforced by simulating a random walk of fictional energy-carrying particles. The use of Monte Carlo simulations for light transport was first proposed in the early '80s [52], and has been continuously developed since then to increase its performances and adapt it to ever more complex geometries [53–55]. Indeed, due to its stochastic nature, the computational burden needed to solve the RTE with the Monte Carlo method is generally high also for simple geometries, but remains basically constant moving to extremely complex meshes and boundary conditions, which can be modeled exactly with little extra effort. Analogously, more realistic phase functions can also be implemented straightforwardly.

One of the main advantages of the MC method is that it lends itself to a very intuitive interpretation in terms of random walks. However, we must keep in mind that this is only

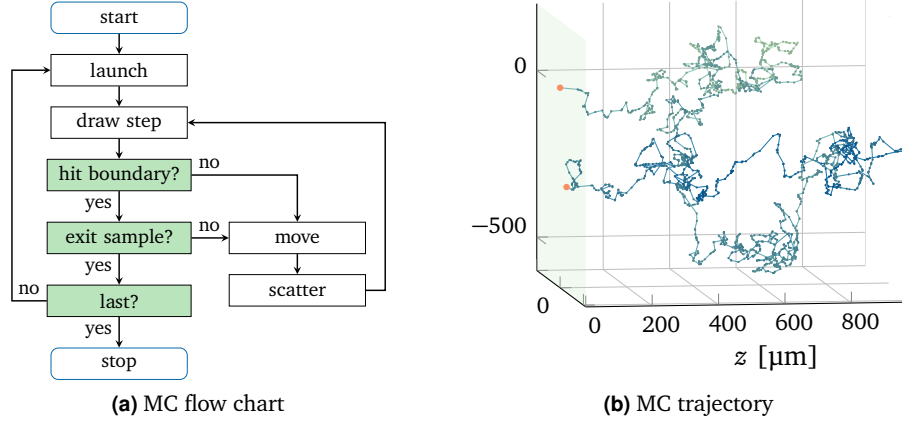


Fig. 2.4.: Simplified flowchart of a Monte Carlo simulation. On the right panel, a typical MC trajectory in a semi-infinite medium ($n = 1.5$) is shown, simulated using a scattering mean free path of $l_s = 10\mu\text{m}$ and a scattering anisotropy of $g = 0.6034$.

a convenient (mathematical) way of solving the RTE and should never be considered as the physical picture. The particles that we refer to are definitely not ‘photons’, and they do not propagate in ‘steps’ nor undergo ‘collisions’. Nonetheless, as we will see also in the following section, linking the problem of (light) transport to the random walk model provides an extremely rich insight at the basis of both the equation of radiative transport and its diffusive approximation.

Building on this random walk analogy, we can express the density of particles propagating along direction \mathbf{s} and relate it to the specific intensity as

$$N(\mathbf{r}, t, \mathbf{s}) = \frac{I(\mathbf{r}, t, \mathbf{s})}{Ev} \quad [\text{m}^{-3} \text{sr}^{-1}] \quad (2.54)$$

where E is the energy per particle. Using this definition, the RTE becomes

$$\frac{1}{v} \frac{\partial N}{\partial t} = -\mathbf{s} \cdot \nabla N - (\mu_s + \mu_a)N + \mu_{\text{tot}} \int_{4\pi} N(\mathbf{r}, t, \mathbf{s}') p(\mathbf{s}, \mathbf{s}') d\Omega' + Q \quad (2.55)$$

where $Q(\mathbf{r}, t, \mathbf{s}) = q(\mathbf{r}, t, \mathbf{s}')/Ev$ is the number of particles emitted per unit time, volume and solid angle. Equation (2.55) can be now solved numerically by tracing particles from the source Q , through the scattering medium. During the propagation, particles are scattered and absorbed with probabilities μ_s and μ_a per unit length, until they are absorbed or exit the sample (see Figure 2.4). In the limit of a large number of repetitions, the MC estimate of radiometric quantities approaches the solution. In practice of course, only a finite number of trajectories is always used, with a number varying heavily depending on the exact geometry, detection scheme and variance reduction techniques applied. The solution obtained is nonetheless *exact*, yet affected by some degree of statistical noise.

The basic principle of Monte Carlo simulations consists of properly sampling probability distributions. Typically, we start by generating a random number ξ uniformly distributed

between 0 and 1, and use it as the building block to obtain a random variable x distributed according to $p(x)$. The target probability distribution $p(x)$ is normalized over the entire domain $a \leq x \leq b$. We can therefore define its cumulative distribution function (CDF)

$$P_x(\tilde{x}) = \int_a^{\tilde{x}} p(x) dx \quad (2.56)$$

which describes the probability of $a \leq x \leq \tilde{x}$, and is therefore also bound between 0 and 1. On the other hand, the cumulative distribution function for the uniformly distributed variable ξ is simply

$$P_\xi(\tilde{\xi}) = \int_0^{\tilde{\xi}} p(\xi) d\xi = \tilde{\xi}. \quad (2.57)$$

To sample $p(x)$, we assume the existence of a nondecreasing function $f(\xi) = x$ mapping $\xi \in [0, 1]$ into $x \in [a, b]$. Its expression is obtained by setting

$$P_x(\tilde{x}) = P_\xi(\tilde{\xi}) \rightarrow \int_a^{\tilde{x}} p(x) dx = \tilde{\xi} \quad (2.58)$$

which, solved for \tilde{x} , yields the sought function $f(\xi)$. A prominent example where this expression admits an analytic solution is that of the exponential distribution, which represents the step size distribution given by the scattering rate μ_s . In fact, we can write the probability of taking a step longer than a certain value $\tilde{\ell}$ as

$$P(\ell \geq \tilde{\ell}) = \exp(-\mu_s \tilde{\ell}) \quad (2.59)$$

which can be rearranged to yield the CDF (2.57)

$$P(\ell < \tilde{\ell}) = \int_0^{\tilde{\ell}} p(\ell) d\ell = 1 - \exp(-\mu_s \tilde{\ell}). \quad (2.60)$$

Solving for $\tilde{\ell}$ gives

$$\ell = \frac{-\ln(1 - \xi)}{\mu_s} \quad (2.61)$$

which is the actual expression used to obtain an exponentially distributed random variate ℓ from a uniformly distributed $\xi \in [0, 1]$ (see also Appendix B). Analogously, in the presence of absorption, the step length to the next inelastic scattering event is given by

$$\ell_i = \frac{-\ln(1 - \xi)}{\mu_a}. \quad (2.62)$$

If $\ell_i < \ell$, the particle reaches the absorption event before the scattering event and is absorbed after the step. Otherwise, the scattering event is reached first and the algorithm proceeds. This way of handling absorption is based on the fact that the exponential distribution, and in particular the absorption process, is *memoryless*. As a matter of fact, it does not matter how far the particle has already traveled, but solely the probability of absorption per unit length traveled μ_a .

When the particle undergoes a scattering event its propagation direction is deflected according to the phase function $p(\mathbf{s}, \mathbf{s}_0)$. As discussed in section 2.1.2, it is commonly assumed that $p(\mathbf{s}, \mathbf{s}_0) = p(\mathbf{s} \cdot \mathbf{s}_0)$ and that the new direction is simply characterized by a pair of deflection and azimuth angles θ and ϕ . The azimuth angle is uniformly distributed in $\phi \in [0, 2\pi)$:

$$\phi = 2\pi\xi. \quad (2.63)$$

A commonly used distribution for the deflection angle is the Henyey-Greenstein distribution (2.30), whose cumulative distribution function can also be inverted to give

$$\cos \theta = \frac{1}{2\pi} \left(1 + g^2 - \frac{1 - g^2}{1 - g + 2g\xi} \right) \quad (2.64)$$

or simply $\cos \theta = 2\xi - 1$ in the case of isotropic scattering ($g = 0$).

Whenever a particle eventually reaches a boundary, the probability to be reflected is usually described by the Fresnel reflection coefficient $R(\theta_i)$ averaged over different polarizations

$$R(\theta_i) = \begin{cases} \frac{1}{2} \left(\frac{n_i \cos \theta_i - n_e \cos \theta_r}{n_i \cos \theta_i + n_e \cos \theta_r} \right)^2 + \frac{1}{2} \left(\frac{n_e \cos \theta_i - n_i \cos \theta_r}{n_e \cos \theta_i + n_i \cos \theta_r} \right)^2 & \text{for } 0 < \theta_i < \theta_c \\ 1 & \text{for } \theta_i \geq \theta_c = \arcsin \frac{n_e}{n_i} \end{cases} \quad (2.65)$$

where θ_i and θ_r are the angles of incidence and refraction calculated using Snell's law

$$\theta_r = \arcsin \left(\frac{n_i}{n_e} \sin \theta_i \right). \quad (2.66)$$

The interaction with the boundary is handled as follows: the particle is moved to the boundary, and a uniformly distributed number $\xi \in [0, 1)$ is compared to $R(\theta_i)$ to decide whether the particle is reflected ($\xi \leq R(\theta_i)$) or transmitted ($\xi > R(\theta_i)$) with an angle θ_r . In both cases, the exceeding part of the step can be either stored and reused to walk the original step length, or it can be discarded and redrawn, taking advantage of the memoryless nature of the exponential distribution.

A prominent feature of the Monte Carlo simulation scheme is that it can naturally accommodate a time-domain simulation, since the total path length traveled by each particle can be readily converted to time dividing by the current velocity of the particle, v . Working in the time-domain offers another relevant advantage, as expressed by the properties of the RTE described at the end of section 2.1.5. On one hand, absorption can be ignored since its effect can be added rigorously to an absorption-free simulation using Lambert-Beer law. Secondly, in any simulation performed with $\mu_a = 0$, the scattering coefficient μ_s can be rescaled by rescaling the spatial and temporal coordinates. While this is particularly useful in certain simple geometries lacking spatially dependent features (such as infinite and semi-infinite media), rescaling can still be profitably applied to the slab geometry, as will be discussed in Chapter 5.

2.2 From Radiative transfer to Diffusion

2.2.1 The diffusion approximation

As we have seen in the previous section, the RTE is a complex equation requiring a significant computational effort to be solved, with no general analytic (closed-form) solution available. Nevertheless, when the transport phenomenon is dominated by multiple scattering, the diffusive approximation emerges as a widely and successfully used model to yield a variety of simple solutions for both steady-state and time-dependent excitation [44, 56].

The first appearance of the concept of diffusion dates back to almost two centuries ago, when J. Fourier published the *Théorie analytique de la Chaleur* (1822), which later on inspired G. Ohm to apply it to the problem of charge conduction and A. Fick who used it to describe the concentration of salts in a solvent. Both Fourier and Fick derived two relations linking the flux with the change of concentration, and the increase of concentration with time due to this flux. These equations have since then been used with great success to predict the transport of heat, particles, mass, charge, population and, as in our case, light in highly turbid media. In this section we will discuss how the diffusion equation (DE) can be derived from the RTE [7, 23], with special emphasis on the approximations required.

The first useful relation to derive the DE is the continuity equation for the energy density, which can be obtained by the RTE (2.51) without the need of any other approximation, just by integrating over the whole solid angle

$$\begin{aligned} \frac{1}{v} \frac{\partial}{\partial t} \int_{\Omega} I(\mathbf{r}, t, \mathbf{s}) d\Omega + \nabla \cdot \int_{\Omega} \mathbf{s} I(\mathbf{r}, t, \mathbf{s}) d\Omega + \mu_{\text{tot}} \int_{\Omega} I(\mathbf{r}, t, \mathbf{s}) d\Omega \\ - \mu_{\text{tot}} \int_{\Omega, \Omega'} I(\mathbf{r}, t, \mathbf{s}') p(\mathbf{s}, \mathbf{s}') d\Omega' d\Omega = \int_{\Omega} q(\mathbf{r}, t, \mathbf{s}) d\Omega \end{aligned} \quad (2.67)$$

where

$$Q(\mathbf{r}, t) = \int_{\Omega} q(\mathbf{r}, t, \mathbf{s}) d\Omega \quad [\text{W m}^{-3}] \quad (2.68)$$

is the source energy density. If we identify the term

$$\mu_{\text{tot}} \int_{\Omega, \Omega'} I(\mathbf{r}, t, \mathbf{s}') p(\mathbf{s}, \mathbf{s}') d\Omega' d\Omega = \mu_{\text{tot}} U(\mathbf{r}, t) \int_{\Omega'} p(\mathbf{s}, \mathbf{s}') d\Omega' = \mu_s U(\mathbf{r}, t) \quad (2.69)$$

using the normalization of the phase function (2.28) and the definition for the average intensity (2.44) and and flux density (2.46), we can write

$$\begin{aligned} \frac{1}{v} \frac{\partial}{\partial t} U(\mathbf{r}, t) + \nabla \cdot \mathbf{F}(\mathbf{r}, t) + \mu_{\text{tot}} U(\mathbf{r}, t) - \mu_s U(\mathbf{r}, t) = Q(\mathbf{r}, t) \\ \rightarrow \frac{1}{v} \frac{\partial}{\partial t} U(\mathbf{r}, t) + \nabla \cdot \mathbf{F}(\mathbf{r}, t) + \mu_a U(\mathbf{r}, t) = Q(\mathbf{r}, t) \end{aligned} \quad (2.70)$$

where eventually the total energy is of course independent of the scattering strength.

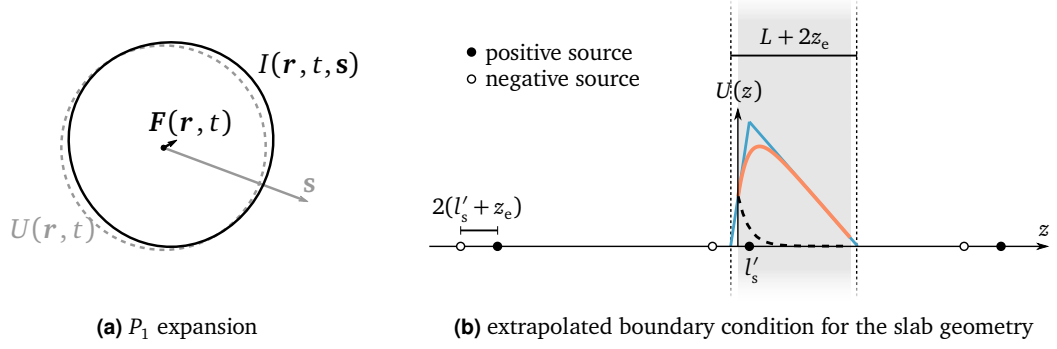


Fig. 2.5.: Panel (a) depicts qualitatively the contribution of the first (isotropic) and second term of the P_1 expansion of the specific intensity, leading to the diffusive approximation. Panel (b) illustrates the procedure used to solve light transport in an infinite plane-parallel scattering medium, using the method of mirror images (see Section 2.2.3). The propagator for the average energy is plotted inside a medium with $L/l'_s = 10$ in blue for a point source placed in $z = l'_s$ and in orange for an array of sources with exponentially decreasing intensity (black dashed line), mimicking the effect of an incoming pencil beam. Extrapolating $U(z)$ along its derivative at the boundaries defines the extrapolated length z_e (see Section 2.2.3).

In the general case of time-dependent sources, obtaining the diffusive equation requires that we apply two simplifying assumptions. The first one assumes that the radiance inside a diffusive medium is almost isotropic, with a slightly unbalanced flux towards \mathbf{s}_j . The simplest approximation that can be made for the diffuse intensity $I(\mathbf{r}, t, \mathbf{s})$ is to assume a series expansion in spherical harmonics

$$I(\mathbf{r}, t, \mathbf{s}) \simeq f_0(\mathbf{r}, t) + f_1(\mathbf{r}, t) \mathbf{s}_j \cdot \mathbf{s}, \quad (2.71)$$

where the expression for f_0 and f_1 can be readily found looking at the definitions (2.44) and (2.46) for the energy and flux densities (see also Appendix A.2), yielding

$$I(\mathbf{r}, t, \mathbf{s}) \simeq \frac{1}{4\pi} U(\mathbf{r}, t) + \frac{3}{4\pi} \mathbf{F}(\mathbf{r}, t) \cdot \mathbf{s}. \quad (2.72)$$

Equation (2.72) is a good approximation for the specific intensity if the contribution of higher-order spherical harmonics is negligible, which is usually verified when $U(\mathbf{r}, t) \gg 3\mathbf{F}(\mathbf{r}, t) \cdot \mathbf{s}$ (see Figure 2.5a).

Now we can use the spherical harmonic expansion (2.72) to simplify the relation that we find multiplying the RTE by \mathbf{s} and then integrating over the solid angle

$$\begin{aligned} \frac{1}{v} \frac{\partial}{\partial t} \mathbf{F}(\mathbf{r}, t) + \int_{\Omega} \left(\mathbf{s} \cdot \nabla \left[\frac{1}{4\pi} U(\mathbf{r}, t) + \frac{3}{4\pi} \mathbf{F}(\mathbf{r}, t) \cdot \mathbf{s} \right] \right) \mathbf{s} d\Omega + \mu_{\text{tot}} \mathbf{F}(\mathbf{r}, t) \\ - \mu_{\text{tot}} \int_{\Omega, \Omega'} p(\mathbf{s}, \mathbf{s}') \left[\frac{1}{4\pi} U(\mathbf{r}, t) + \frac{3}{4\pi} \mathbf{F}(\mathbf{r}, t) \cdot \mathbf{s} \right] \mathbf{s} d\Omega' d\Omega = \int_{\Omega} q(\mathbf{r}, t, \mathbf{s}) \mathbf{s} d\Omega. \end{aligned} \quad (2.73)$$

After some manipulation, this expression can be rewritten as (see Appendix A.2)

$$\left(\frac{1}{v} \frac{\partial}{\partial t} + (\mu_a + \mu'_s)\right) \mathbf{F}(\mathbf{r}, t) + \frac{\nabla U(\mathbf{r}, t)}{3} = 0 \quad (2.74)$$

assuming an isotropic source $q(\mathbf{r}, t, \mathbf{s}) = q(\mathbf{r}, t)/4\pi$.

At this stage, the second simplifying assumption leading to the diffusive equation is that of neglecting the time variation of the diffuse flux $\mathbf{F}(\mathbf{r}, t)$ over a time range l'_s/v , assuming that it is much smaller than the vector itself

$$\frac{1}{\mu'_s v} \left| \frac{\partial \mathbf{F}(\mathbf{r}, t)}{\partial t} \right| \ll |\mathbf{F}(\mathbf{r}, t)|. \quad (2.75)$$

It is worth commenting further this approximation, since by removing the temporal dependence of $\mathbf{F}(\mathbf{r}, t)$ we are effectively invalidating one of the fundamental similarity relations of the RTE, relating the specific intensity in the presence of absorption with the specific intensity in a non-absorbing medium (2.53). This issue has been thoroughly discussed by several authors [7, 57–59] eventually deriving an expression that varies depending on whether the experimental measurements are performed in the spatial or in the temporal domain. Indeed, the diffusion approximation is expected to hold for light that has undergone a multitude of scattering events, and therefore absorption obstructs the diffusive regime in that it selectively extinguishes the light that could propagate into a deeply multiple scattering regime. In a steady state detection scheme, the acquired signal would be dominated by low-order scattering and be poorly modeled by the diffusive approximation. Turning to a richer time-domain representation of transport, however, the problem becomes easier to handle since by addressing separately different times we avoid mixing the contributions of low-order and high-order scattered light. In this framework, which is the one relevant for this work, it is more appropriate to derive the diffusive approximation considering a non-absorbing medium, and then adding the effect of absorption separately, in a way that preserves the original symmetry of the RTE [7, 44, 60].

By doing so, we can rewrite the two expressions (2.70) and (2.74) as

$$\frac{1}{v} \frac{\partial U(\mathbf{r}, t)}{\partial t} + \nabla \mathbf{F}(\mathbf{r}, t) = Q(\mathbf{r}, t) \quad (2.76)$$

$$\mathbf{F}(\mathbf{r}, t) = -\frac{1}{v} D \nabla U(\mathbf{r}, t) \quad (2.77)$$

where we have introduced a diffusion coefficient

$$D = \frac{v}{3\mu'_s} \quad (2.78)$$

independent of absorption, which allows us to treat the dependence on absorption in accordance with the RTE without the additional approximations intrinsic to the diffusion framework.

The diffusion equation for a homogeneous medium is finally obtained by substituting equation (2.77) into (2.76), yielding

$$\frac{1}{v} \left(\frac{\partial}{\partial t} - D \nabla^2 \right) U(\mathbf{r}, t) = Q(\mathbf{r}, t). \quad (2.79)$$

whose solution, assuming a simple point source $Q(\mathbf{r}, t) = E_0 \delta(\mathbf{r}) \delta(t)$, is given by an expanding Gaussian profile

$$U(\mathbf{r}, t) = \frac{v E_0}{(4\pi D t)^{3/2}} \exp\left(-\frac{r^2}{4Dt}\right) \quad (2.80)$$

for an unbounded medium.

2.2.2 Diffusion reloaded: the random walk picture

Before we sum up the approximations that we used to derive the diffusive equation and its validity range, it is useful to review an alternative derivation based on the random walk model. To this purpose, let us consider an ensemble of random walkers taking steps of random length ℓ in d dimensions, assuming that the direction of each step is completely uncorrelated with the previous one. The position \mathbf{r}_n that a walker will reach after n steps will be given by

$$\mathbf{r}_n = \sum_{i=1}^n \mathbf{x}_i = \sum_{i=1}^n \ell_i \mathbf{x}_i \quad (2.81)$$

where $\mathbf{x}_i = \ell_i \mathbf{x}_i$ is a set of random identically distributed variables according to some function $p(\mathbf{x}) = p(|\mathbf{x}|)$ which depends only on the magnitude of the step length. For symmetry reasons, in the limit of many repetitions, the expectation value for the sum (2.81) will still be centered in $\langle \mathbf{r} \rangle = 0$. Conversely, the expectation value of the mean square displacement (MSD) will be

$$\langle \mathbf{r}_n^2 \rangle = \left\langle \sum_{i=1}^n \mathbf{x}_i \cdot \sum_{j=1}^n \mathbf{x}_j \right\rangle = \sum_{i=1}^n \langle \mathbf{x}_i^2 \rangle + \left\langle \sum_{i=1}^n \sum_{j \neq i}^n \mathbf{x}_i \cdot \mathbf{x}_j \right\rangle = n \langle \ell^2 \rangle, \quad (2.82)$$

where the cross terms vanishes on average due to the absence of correlations between steps. If we introduce the probability density function (PDF) $P_n(\mathbf{r})$ associated with the final position \mathbf{r} after n steps, we can define the following recursion relation

$$P_{n+1}(\mathbf{r}) = \int p(\mathbf{x}) P_n(\mathbf{r} - \mathbf{x}) d^d \mathbf{x} \quad (2.83)$$

exploiting the fact that for uncorrelated steps we can factorize the integrand. Equation (2.83), which in $d = 1$ is sometimes referred to as Bachelier's equation, can be seen as a continuity equation for the total number of random walkers: the probability of being in \mathbf{r} after $n + 1$ steps is equal to the probability of being at $\mathbf{r} - \mathbf{x}$ at the previous step, multiplied by the probability of reaching \mathbf{r} in one step. In the limit of a large number of

steps, we can assume that $P_n(\mathbf{r})$ will be a rather smooth distribution, with significant variations over a length scale much longer than the typical step size. This allows us to expand it around \mathbf{r} as

$$\begin{aligned} P_{n+1}(\mathbf{r}) &= \int p(\mathbf{x}) \left[P_n(\mathbf{r}) - \mathbf{x} \cdot \nabla P_n(\mathbf{r}) + \frac{1}{2} \mathbf{x} \cdot \nabla \nabla P_n(\mathbf{r}) \cdot \mathbf{x} + \dots \right] d^d \mathbf{x} \\ &= P_n(\mathbf{r}) + \frac{\langle \ell^2 \rangle}{2d} \nabla^2 P_n(\mathbf{r}) + \dots \quad (2.84) \end{aligned}$$

where only the even terms will survive integration. If we limit to the first two terms of the expansion and divide by the average time $\langle \Delta t \rangle = \langle \ell \rangle / v$ needed to complete one step we obtain

$$\frac{P_{n+1}(\mathbf{r}) - P_n(\mathbf{r})}{\Delta t} = \frac{\langle \ell^2 \rangle}{2d \Delta t} \nabla^2 P_n(\mathbf{r}). \quad (2.85)$$

In the limit $n \rightarrow \infty$ the distribution $P_n(\mathbf{r}) \rightarrow \psi(\mathbf{r}, t = n\Delta t)$ satisfies a diffusion equation

$$\frac{\partial \psi(\mathbf{r}, t)}{\partial t} = D \nabla^2 \psi(\mathbf{r}, t) \quad (2.86)$$

whose solution is the same of equation (2.80)

$$\psi(\mathbf{r}, t) = \frac{e^{-r^2/4Dt}}{(4\pi Dt)^{d/2}} \quad (2.87)$$

with

$$D = \frac{\langle \ell^2 \rangle}{2d \Delta t} = \frac{v}{2d} \frac{\langle \ell^2 \rangle}{\langle \ell \rangle} \quad (2.88)$$

as the diffusion coefficient of the random walk process. If we consider, as in our Monte Carlo simulations, an exponential step length distribution

$$p(\ell) = \mu_s e^{-\mu_s \ell} \quad (2.89)$$

which has $\langle \ell \rangle = \mu_s^{-1}$ and $\langle \ell^2 \rangle = 2\mu_s^{-2}$ we recover exactly the same expression (2.78) that we have found in $d = 3$ with the P_1 approximation of the RTE.

This alternative derivation of the diffusion equation reveals valuable information on the approximations and simplifications that are not obvious when taking the conventional approach. Firstly, we can now establish a connection between diffusion and the Central limit Theorem, which states that the sum of a large number of independent random variables with finite first and second moments tends to a Gaussian distribution. In this respect, the diffusion equation is just a representation of the same principle. In both cases, the first requirement is that the number of steps must be $n \gg 1$. Secondly, this alternative derivation allows us to shed light on the implications of the P_1 approximation, which, in the random walk model, amounts to neglecting moments of the step length distribution higher than the third (equation (2.84)). In this framework, we also obtain

a measure of the breakdown of the diffusive approximation by evaluating the excess kurtosis

$$\gamma_2 = \frac{\langle \ell^4 \rangle}{\langle \ell^2 \rangle^2} - 3 \quad (2.90)$$

of the time-resolved density/intensity distributions [60]. Finally, the random walk derivation gives us some complementary insight over the role of absorption. The fundamental requirement for diffusion theory is that the number of step must be large, or equivalently that only late times can be considered. Absorption attenuates late light and should be therefore be preferably small, but this is just a practical limitation dictated by the sensitivity of an experiment, rather than by diffusion theory itself. This is confirmed by the very definition of the diffusion coefficient, which describes how the variance of the spatial distribution grows in time. In the limit of a large number of steps (i.e., the discrete and continuous time domain are equivalent), adding absorption does not change the shape of the distribution since at any given time the particles will have walked the same distance. It is therefore clear that the variance of the spatial distribution, and therefore the diffusion coefficient, is independent of absorption.

2.2.3 Diffusion in bounded media

Light propagation in finite geometries introduces the problem of appropriately handling boundaries between different media, where the almost-isotropic assumption of the P_1 approximation is violated. For a diffusive medium bounded by a convex or flat surface Σ at the interface with a non-scattering region, the exact boundary condition for the radiance $I(\mathbf{r}, t, \mathbf{s})$ is that there should be no diffuse light entering the medium from outside through the interface Σ . Any intensity at $\mathbf{r} \in \Sigma$ coming from a direction \mathbf{s} directed towards the diffusive medium can only originate from reflection at the boundary

$$I(\mathbf{r}, t, \mathbf{s}) = R(\mathbf{s}' \cdot \mathbf{q} = \cos \theta_i) I(\mathbf{r}, t, \mathbf{s}') \quad (2.91)$$

where \mathbf{q} is the unit vector normal to Σ and $R(\cos \theta_i)$ is the Fresnel reflection coefficient for unpolarized light (2.65). With the simple angular distribution assumed in the expansion (2.72), the requirement (2.91) cannot be satisfied exactly, and approximate boundary conditions must be considered. It is assumed that the condition (2.91) is verified on average for all inward pointing directions \mathbf{s}

$$-\int_{\mathbf{s} \cdot \mathbf{q} < 0} I(\mathbf{r}, t, \mathbf{s})(\mathbf{s} \cdot \mathbf{q}) d\Omega = \int_{\mathbf{s} \cdot \mathbf{q} > 0} R(\theta_i) I(\mathbf{r}, t, \mathbf{s})(\mathbf{s} \cdot \mathbf{q}) d\Omega \quad (2.92)$$

which is a boundary condition for the total radiation coming from the boundary surface. Making use of the angular distribution for the specific intensity (2.72) and calculating the integrals of equation (2.92), the boundary condition for the fluence rate can be written as (see Appendix A.3)

$$[U(\mathbf{r}, t) - 2AF(\mathbf{r}, t) \cdot \mathbf{q}]_{\mathbf{r} \in \Sigma} = 0 \quad (2.93)$$

with

$$A = \frac{1 + 3 \int_0^{\pi/2} R(\theta_i) \cos^2 \theta_i \sin \theta_i d\theta_i}{1 - 2 \int_0^{\pi/2} R(\theta_i) \cos \theta_i \sin \theta_i d\theta_i} \quad (2.94)$$

representing a coefficient that depends only on the relative refractive index $n = n_i/n_e$, which is $A = 1$ if $n_i = n_e$ and $A > 1$ otherwise. This boundary condition is denoted as the partial current boundary condition (PCBC) and represents the most accurate boundary condition for light diffusion at a boundary. The PCBC can be recast differently in terms of the fluence alone using Fick's law (2.77) to write

$$\left[U(\mathbf{r}, t) - \frac{2AD}{v} \frac{\partial}{\partial q} U(\mathbf{r}, t) \right]_{\mathbf{r} \in \Sigma} = 0. \quad (2.95)$$

According to this condition (which is sometimes called Robin boundary condition), the derivative of $U(\mathbf{r}, t)$ along the direction normal to the boundary is proportional to $U(\mathbf{r}, t)$ itself. An extrapolated decrement of $U(\mathbf{r}, t)$ inside the non-scattering region is obtained if the derivative of $U(\mathbf{r}, t)$ is assumed to remain constant in the non-scattering region to the value on the boundary (see Figure 2.5b). The distance from the geometrical boundary at which $U(\mathbf{r}, t)$ is extrapolated to zero is denoted as the extrapolated distance z_e

$$z_e = \frac{2AD}{v} = \frac{2}{3} Al'_s. \quad (2.96)$$

The boundary condition that assumes $U = 0$ on the surface at the extrapolated distance z_e is denoted as the extrapolated boundary condition (EBC) [44, 61].

When describing light propagation in bounded media, few words must be spent on how to model the source term, which will be usually placed outside the turbid medium. Moreover, so far we have always assumed perfectly isotropic point sources, which are very far from usual experimental conditions. If we consider a collimated beam of light incident on a scattering medium or being delivered through an optical fiber, in each of these cases light starts propagating with a well defined initial direction. The most common way of modeling this situation is to approximate the directed source with an isotropic point source *inside* the medium, placed at a depth $z_{\text{src}} = l'_s$ along the direction of the beam (see Figure 2.5b). This is motivated by the fact that z_{src} represents the mean depth at which the first (isotropic-equivalent) scattering event occurs in the case of the exponential distribution [56]. This approximation is of course acceptable only if we perform our measurements very far (both in time and space) from the source. A better approximation is that of modeling the directed source as a distribution of isotropic sources with intensities proportional to $\exp(-z\mu'_s)$, where z is the distance from the interface.

One of the most relevant bounded configuration that has been widely studied in the past decades is the infinitely extended plane-parallel slab. Indeed, an array of different physical systems are often represented as a layer or a combination of layers, ranging from atmospheric physics [4] to geosciences [8], paint and coatings applications [62] and biological tissues [6], to name a few. For this reason, the solution of the diffusion

problem in the slab geometry is one of the most relevant and widely used in light transport studies.

To obtain the solution to this particular problem we follow the method of mirror images [44, 56]. The method involves the use of boundary conditions that assume a vanishing fluence at some distance from the physical boundary (e.g., at $z_e = 2AD$ for the EBC). For geometries with regular boundaries such as the slab, the method of images allows to assemble the solution for the fluence inside the medium as a superposition of (infinite) solutions for the infinite medium. As a matter of fact, the series converges extremely quickly, and a few terms are sufficient in almost all applications of interest. In a sense, the EBC can be seen as a mapping from a sample of thickness L to a sample with an effective thickness $L_{\text{eff}} = L + 2z_e$ such that the new effective sample can be considered as infinitely extended (i.e., fluence goes to 0 at its ‘effective’ boundaries). The flux exiting the diffusive medium, representing in this case the reflectance and transmittance from the slab, is obtained by applying Fick’s law at the boundary of the medium.

An hybrid heuristic approach, based both on the EBC and the PCBC (also sometimes named extrapolated boundary partial current (EBPC) [63]) has also emerged more recently in the literature [64], and will be briefly reviewed at the end of the subsection.

The geometry of the problem with a description of the notation used is shown in Figure 2.5b. We consider an isotropic point source of unit strength $q(\mathbf{r}, t) = \delta^3(\mathbf{r} - \mathbf{r}_{\text{src}})\delta(t)$ placed at $\mathbf{r}_{\text{src}} = (0, 0, l'_s)$ and $L > l'_s$ thickness of the slab. According to the EBC, the fluence is assumed equal to 0 at two extrapolated flat surfaces outside the turbid medium at the extrapolated distance z_e from the physical boundaries of the slab. This condition is enforced by the method of images by using, in addition to the real source in \mathbf{r}_{src} , an infinite number of pairs of positive and negative sources in an infinite diffusive medium having the same optical properties of the slab. The locations \mathbf{r}_m^\pm of the first few positive and negative sources is shown in Figure 2.5b, and is such that the fluence of each source is balanced by an image source of opposite sign placed at a symmetric position with respect to both extrapolated surfaces. The only real source is placed at $z_0^+ = l'_s$. All other sources are image sources, and are placed along the z -axis at

$$\begin{cases} z_m^+ &= 2m(L + 2z_e) + l'_s \\ z_m^- &= 2m(L + 2z_e) - 2z_e - l'_s \end{cases}$$

for $m = \pm 1, \pm 2, \dots \pm \infty$. Adding the contributions of all the source pairs, the Green’s function for the fluence rate at $\mathbf{r} = (x, y, z)$ results in

$$U(\mathbf{r}, t) = \frac{\nu \exp\left(-\frac{\rho^2}{4Dt} - \mu_a \nu t\right)}{(4\pi Dt)^{3/2}} \sum_{m=-\infty}^{\infty} \left\{ \exp\left[-\frac{(z - z_m^+)^2}{4Dt}\right] - \exp\left[-\frac{(z - z_m^-)^2}{4Dt}\right] \right\} \quad (2.97)$$

for $0 \leq z \leq L$. The flux and the specific intensity can be calculated using (2.77) and (2.72), allowing to retrieve the time-resolved transmittance and reflectance as

$$T(\rho, t) = -\frac{D}{\nu} \frac{\partial}{\partial z} U(\rho, z = L, t) \quad [\text{W m}^{-1}] \quad (2.98)$$

$$R(\rho, t) = \frac{D}{v} \frac{\partial}{\partial z} U(\rho, z = 0, t). \quad [\text{W m}^{-1}] \quad (2.99)$$

In a random walk picture, given the fact that the functions $R(\rho, t)$ and $T(\rho, t)$ correspond to a unit energy point-like source, they can also be considered as the probability per unit of time and area that a walker emitted at \mathbf{r}_{src} at $t = 0$, exits at time t at a distance ρ from the z -axis. Making use of equation (2.97) we obtain

$$T(\rho, t) = \frac{\exp\left(-\frac{\rho^2}{4Dt} - \mu_a v t\right)}{2(4\pi D)^{3/2} t^{5/2}} \sum_{m=-\infty}^{\infty} \left\{ z_{1,m} \exp\left[-\frac{z_{1,m}^2}{4Dt}\right] - z_{2,m} \exp\left[-\frac{z_{2,m}^2}{4Dt}\right] \right\} \quad (2.100)$$

$$R(\rho, t) = \frac{\exp\left(-\frac{\rho^2}{4Dt} - \mu_a v t\right)}{2(4\pi D)^{3/2} t^{5/2}} \sum_{m=-\infty}^{\infty} \left\{ z_{3,m} \exp\left[-\frac{z_{3,m}^2}{4Dt}\right] - z_{4,m} \exp\left[-\frac{z_{4,m}^2}{4Dt}\right] \right\} \quad (2.101)$$

with

$$\begin{cases} z_{1,m} &= (1-2m)L - 4mz_e - l'_s \\ z_{2,m} &= (1-2m)L - (4m-2)z_e + l'_s \\ z_{3,m} &= -2mL - 4mz_e - l'_s \\ z_{4,m} &= -2mL - (4m-2)z_e + l'_s \end{cases} \quad (2.102)$$

and $\rho = \sqrt{x^2 + y^2}$. Equations (2.100) and (2.101) are infinite series and should be truncated for practical applications. Since the distance of the mirror sources from the boundaries increases with increasing m , the contribution of high- m sources is expected to be significant only of large values of ρ and/or t . Figures 2.6a and 2.6b show the transmittance and reflectance profiles, respectively, calculated at different times for a non-absorbing slab sample in air with $L = 10$ mm, $l'_s = l_s = 0.1$ mm and $n_{\text{in}} = 1.5$. Data simulated with the Monte Carlo method relative to a sample with anisotropic scattering ($g = 0.6034$) are plotted together for comparison, showing the good agreement with the theory, the validity of the similarity relation (2.38) and of the point source approximation (in the simulation, a pencil beam source is used). Figure 2.7 shows how the transmittance $T(\rho, t)$ (the case for $R(\rho, t)$ is analogous) depends on various parameters at a fixed delay. According to the diffusive approximation, most parameters, including the thickness L and the refractive index contrast n , only affect the amplitude of the time-resolved profile, while only a change in l'_s modifies its shape.

This is reflected in the prediction cast by diffusion theory for the temporal expansion of the mean square width of $T(\rho, t)$ (and analogously for $R(\rho, t)$), defined as

$$w^2(t) = \frac{\int_0^\infty \rho^2 T(\rho, t) \rho \, d\rho}{\int_0^\infty T(\rho, t) \rho \, d\rho} = \frac{\int_0^\infty \rho^2 T(\rho, t) \rho \, d\rho}{T(t)} \quad (2.103)$$

which, for a Gaussian profile with standard deviation σ , is simply given by

$$w^2(t) = 2\sigma^2 = 4Dt. \quad (2.104)$$




Fig. 2.6.: (a)–(b) Spatially-resolved transmittance and reflectance at different times as calculated for a homogeneous scattering slab with $L = 10$ mm, $\mu_a = 0$, $\mu'_s = 10 \text{ mm}^{-1}$, $z_{\text{src}} = 1/\mu'_s$ and $n_{\text{in}} = 1.5$. At each instant the profiles for $T(\rho, t)$ and $R(\rho, t)$ are two identical Gaussian profiles (except for their amplitude) with a mean square width linearly increasing with time as $4Dt$. Panel (c) shows the spatially-integrated transmittance and reflectance for the same set of parameters. Note that, in principle, both equations (2.100) and (2.105) can be formally calculated for any $t > 0$ and give a non-null intensity up to $\rho \rightarrow \infty$, which is unphysical given that light propagates at finite speed. The output of a MC simulation for a similar sample with $g = 0.6034$ and $\mu_s = 10/(1 - g) \text{ mm}^{-1}$ ($\mu'_s = 10 \text{ mm}^{-1}$) is shown for comparison.

We should stress, however, that the linear growth of the mean square width (MSW) with time is not limited to the Gaussian profiles predicted by the diffusion approximation for a slab, but lies rather at the very definition of diffusive transport, as we will discuss further in the following chapters. Figure 2.8 shows a comparison between the diffusion approximation and the output of a Monte Carlo simulation of a sample satisfying the similarity relation, representing the exact solution of the scalar RTE, showing excellent agreement both for the linear MSW growth and the vanishing excess kurtosis γ_2 (2.90).

By integrating equations (2.100) and (2.101) over the infinitely extended exit surfaces, the total time-resolved transmittance and reflectance are obtained as

$$T(t) = \frac{\exp(-\mu_a vt)}{2(4\pi D)^{1/2} t^{3/2}} \sum_{m=-\infty}^{\infty} \left\{ z_{1,m} \exp\left[-\frac{z_{1,m}^2}{4Dt}\right] - z_{2,m} \exp\left[-\frac{z_{2,m}^2}{4Dt}\right] \right\} \quad (2.105)$$

$$R(t) = -\frac{\exp(-\mu_a vt)}{2(4\pi D)^{1/2} t^{3/2}} \sum_{m=-\infty}^{\infty} \left\{ z_{3,m} \exp\left[-\frac{z_{3,m}^2}{4Dt}\right] - z_{4,m} \exp\left[-\frac{z_{4,m}^2}{4Dt}\right] \right\} \quad (2.106)$$

and are plotted in Figure 2.6c for the same illustrative sample. The functions (2.105) and (2.106) represent the Green's functions for an infinitely extended detector. For the reciprocity principle, the functions $T(t)$ and $R(t)$ can be used to describe the time-resolved transmittance and reflectance when an infinitely wide beam with constant radiance impinges perpendicularly on the surface of the slab. At late times, the two curves tend to the same value, meaning that the walkers eventually have the same probability to leave the sample from either side of the slab. This happens because at late




Fig. 2.7.: Dependence of the transmittance $T(\rho, t)$ in the diffusive approximation on (a) the thickness of the slab L , (b) the reduced scattering mean free path l'_s , (c) the refractive index contrast n and (d) the absorption length l_a . In the diffusive approximation, the only parameter that affects the shape of the profile is $l'_s \propto D$, while all others modify only its amplitude.

times the energy density inside the slab tends to a spatial distribution that is symmetric with respect to the middle of the slab.

Perhaps the most characterizing feature of the spatially-integrated transmittance and reflectance is their asymptotic decay. The time constant τ of this exponential represents an important time-scale associated with the transport process. Its definition in terms of the optical parameters becomes apparent by recasting equation (2.105) and (2.106) using the Poisson summation rule [65] to obtain

$$T(t) = -\frac{2\pi D}{L_{\text{eff}}} \sum_{m=1}^{\infty} m \sin\left(\frac{m\pi(l'_s + z_e)}{L_{\text{eff}}}\right) \cos\left(\frac{m\pi(L + z_e)}{L_{\text{eff}}}\right) \exp\left(-\frac{m^2\pi^2 D t}{L_{\text{eff}}^2} - \mu_a v t\right) \quad (2.107)$$

$$R(t) = \frac{2\pi D}{L_{\text{eff}}} \sum_{m=1}^{\infty} m \sin\left(\frac{m\pi(l'_s + z_e)}{L_{\text{eff}}}\right) \cos\left(\frac{m\pi z_e}{L_{\text{eff}}}\right) \exp\left(-\frac{m^2\pi^2 D t}{L_{\text{eff}}^2} - \mu_a v t\right) \quad (2.108)$$

where it can be seen that, at late times, only the term with $m = 1$ survives since all other terms decay exponentially faster. The decay constant associated with such term is the asymptotic decay time of $T(t)$ and $R(t)$ and is given by

$$\frac{1}{\tau} = \frac{\pi^2 D}{(L + 2z_e)^2} + \mu_a v. \quad (2.109)$$

Figure 2.9 shows how $T(t)$ (and its asymptotic decay time τ) depends on the optical and geometric parameters of our example slab. As can be seen, the spatially integrated fluxes have a more complex dependence with time, and all parameters can affect both their shape and amplitude in similar ways, which makes it more difficult to use these functions alone to retrieve more than one such parameters at once.




Fig. 2.8.: (a) Comparison between the MSW of the $T(\rho, t)$ and $R(\rho, t)$ profiles, growing linearly as $4Dt$, and the output of a MC simulation. (b) Both transmittance and reflectance profiles become quickly Gaussian, as shown by the vanishing value of γ_2 . (c) Comparison between the (steady-state) angular distribution predicted by the diffusive approximation $F(\cos \theta_e)$ (see subsection 2.2.4) and the output of the MC simulation. The angular distribution for a Lambertian surface ($F = 1/\pi$) is plotted as a dashed line for comparison. The simulated angular distribution for reflected light exhibits a slight deviation from the theory, due to the proximity to the (pencil beam) source. Deviations are progressively reduced as early-reflected light is rejected from integration.

Throughout this thesis work, a hybrid heuristic approach has been followed to calculate the outgoing time-dependent fluxes, based on the more recently developed EBPC condition. Following this method, the outgoing flux $F_{\text{out}}(\mathbf{r}, t)$ is obtained simply by applying the PCBC relation (2.93) to the solution for the fluence obtained using the EBC (2.97), resulting in the expression

$$F_{\text{out}}(\mathbf{r}, t) = \frac{U(\mathbf{r}, t)}{2A} \quad (2.110)$$

which can be applied to any point of the boundary. Since the PCBC is expected to be less approximated than Fick's law, equation (2.110) should provide more accurate results than the previously derived equations. As a matter of fact, the discontinuity in the optical properties occurring at a boundary may determine strong variations of the flux that might depart, near the physical boundaries, from the assumptions (2.72) and (2.75) needed to obtain the diffusion equation from the RTE. Consequently, Fick's law (2.77) might be compromised near the boundaries. Equation (2.110) is also undoubtedly simpler to use, since the same expression is applicable to any point of the boundary. As an example, for the slab, the same expression (2.110) can be used to evaluate both the time-resolved transmittance and reflectance writing

$$T(\rho, t) = \frac{U(\rho, z = L, t)}{2A} \quad (2.111)$$

$$R(\rho, t) = \frac{U(\rho, z = 0, t)}{2A} \quad (2.112)$$

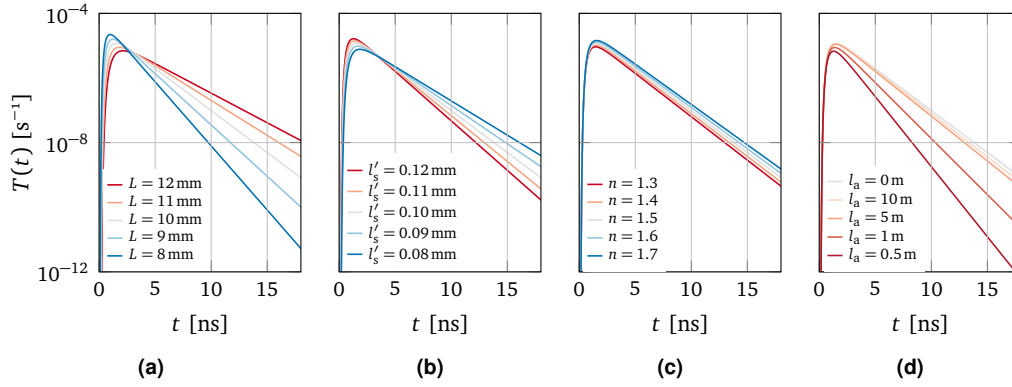


Fig. 2.9.: Dependence of the integrated transmittance $T(t)$ in the diffusive approximation on (a) the thickness of the slab L , (b) the reduced scattering mean free path l'_s , (c) the refractive index contrast n and (d) the absorption length l_a . Each parameter contributes differently to change both the amplitude and the shape of the profile.

and the spatially integrated fluxes

$$T(t) = \frac{U(z = L, t)}{2A} \quad (2.113)$$

$$R(t) = \frac{U(z = 0, t)}{2A}, \quad (2.114)$$

using the definition (2.97) and

$$U(z, t) = \int_0^\infty U(r, t) 2\pi\rho \, d\rho = \frac{ve^{-\mu_a vt}}{\sqrt{4\pi Dt}} \sum_{m=-\infty}^{\infty} \left\{ \exp\left[-\frac{(z-z_m^+)^2}{4Dt}\right] - \exp\left[-\frac{(z-z_m^-)^2}{4Dt}\right] \right\} \quad (2.115)$$

respectively. As for the solutions obtained using Fick's law, for a non-absorbing medium the above expressions verify energy conservation between the total transmitted and reflected energy. It should be pointed out that, in spite of the different expressions obtained following the two different approaches to evaluate the outgoing flux, the final results are in many cases, including the ones previously plotted, indistinguishable. Extremely small differences are appreciable at very early times and are more pronounced in the presence of a high refractive index contrast. Compared with MC simulations, the EBPC solution provides a slightly better description of the time-resolved outgoing flux, especially for the reflectance at short distances from the source. Therefore, the solution is preferable for inversion procedures aimed at retrieving the optical properties of the medium from time-resolved measurements, and as such it has been used throughout the rest of the thesis. For all practical situations described in this work, though, the difference between equation (2.105) and (2.113) is completely negligible and using either model would result in the same output. Considering late times only, it can in fact be shown that the ratio between the two solutions converges exactly to 1.

2.2.4 Angular dependence of outgoing radiance

Measurements of diffused light aimed at characterizing turbid media are often based on detection of scattered light emerging from its outer surface. The flux is commonly collected using optical elements and detectors having a limited numerical aperture. Therefore, for many different reasons the quantity actually measured per unit area is the outgoing radiance accepted by the detection apparatus, which can be written as

$$P(t) = \int_{\Omega_d} I_e(\mathbf{r}, \mathbf{s}_e, t) \mathbf{s}_e \cdot \mathbf{q} d\Omega_d, \quad (2.116)$$

where Ω_d is the acceptance solid angle of the detection system, $I_e(\mathbf{r}, \mathbf{s}_e, t)$ is the specific intensity on the external boundary of the medium, \mathbf{q} is the outwardly directed normal and \mathbf{s}_e is the unit vector pointing outside of the medium with $\mathbf{s}_e \cdot \mathbf{q} = \cos \theta_e$. Knowing the angular distribution of the outgoing radiance is key to understanding if the previously reported solutions for the flux are suitable to describe the actually measured quantity. An analytical expression based on the diffusion approximation and the PCBC can be obtained for the angular dependence of the radiance outgoing from a diffusive medium bounded by a non-scattering region [66]. According to the diffusion approximation, the radiance is assumed to be almost isotropic taking only the first two terms of the spherical harmonic expansion (2.72). The radiance on the external surface can be represented as the fraction of internal radiance I_i that is transmitted in the external medium

$$I_e(\mathbf{r}, \mathbf{s}_e, t) = \frac{n_e^2}{n_i^2} [1 - R(\theta_i)] I_i(\mathbf{r}, \mathbf{s}_i, t), \quad (2.117)$$

where $R(\theta_i)$ is the Fresnel coefficient for unpolarized light (2.65), \mathbf{s}_e and \mathbf{s}_i are related by Snell's law and the term $(n_e/n_i)^2$ accounts for the refraction of the solid angle. If we write (see also Figure A.1)

$$I_i(\mathbf{r}, \mathbf{s}_i, t) = \frac{1}{4\pi} U_i(\mathbf{r}, t) + \frac{3}{4\pi} (\mathbf{F}_{q,i} \cos \theta_i + \mathbf{F}_{u,i} \cos \phi_i \sin \theta_i) \quad (2.118)$$

the normal component of the flux $\mathbf{F}_{u,i}$ averages out with the azimuthal angle, and we arrive to the expression

$$I_e(\mathbf{r}, \mathbf{s}_e, t) = \frac{n_e^2}{n_i^2} [1 - R(\theta_i)] \frac{U_i(\mathbf{r}, t)}{4\pi} \left(1 + \frac{3}{2A} \cos \theta_i \right) \quad (2.119)$$

where we have used the PCBC at the interface (2.93) to express everything in terms of $U(\mathbf{r}, t)$. Substituting this expression into equation (2.116) we obtain

$$P(t) = \frac{n_e^2}{n_i^2} \int_{\Omega_d} d\Omega_d \left\{ 1 - R \left[\arcsin \left(\frac{n_e}{n_i} \sin \theta_e \right) \right] \right\} \frac{U_i(\mathbf{r}, t)}{4\pi} \times \left\{ 1 + \frac{3}{2A} \cos \left[\arcsin \left(\frac{n_e}{n_i} \sin \theta_e \right) \right] \right\} \mathbf{s}_e \cdot \mathbf{q}. \quad (2.120)$$

Since integration over the solid semi-angle leads to the total outgoing flux (2.110), equations (2.119) and (2.120) can be rewritten as

$$I_e(\mathbf{r}, \mathbf{s}_e, t) = \frac{U_i(\mathbf{r}, t)}{2A} F(\theta_e) \rightarrow P(t) = \frac{U_i(\mathbf{r}, t)}{2A} \int_{\Omega_d} F(\theta_e) \cos \theta_e \quad (2.121)$$

with

$$F(\theta_e) = \frac{1}{4\pi} \frac{n_e^2}{n_i^2} \left\{ 1 - R \left[\arcsin \left(\frac{n_e}{n_i} \sin \theta_e \right) \right] \right\} \left\{ 2A + 3 \cos \left[\arcsin \left(\frac{n_e}{n_i} \sin \theta_e \right) \right] \right\} \quad (2.122)$$

and

$$\int_{2\pi} F(\theta_e) \cos \theta_e d\Omega_e = 1. \quad (2.123)$$

Therefore, the angular dependence of the outgoing radiance (averaged over the azimuthal angle) is completely separated from the spatial and temporal dependence and is fully described by $F(\theta_e)$. This function is independent of the geometry considered for the boundary and represents the distribution function for the direction of the outgoing radiation, depending only on the refractive index contrast $n = \frac{n_i}{n_e}$. Still, it should be pointed out that equation (2.122) has been obtained using Snell's law, and therefore it is rigorously applicable only if the external surface of the diffusive medium is sufficiently smooth. Figure 2.8c shows the dependence of $F(\cos \theta_e)$ for the test sample with $n = 1.5$. As can be seen, the distribution is quite different from that of a Lambertian surface (considering Lambert's cosine law, i.e., a surface emitting a radiance independent of the direction, for which $F = 1/\pi$). It could be expected that the angular distribution of light predicted using the simple P_1 decomposition is a rough approximation of the actual distribution since near the boundary, due to the discontinuous variation of the optical properties, strong variations of the flux are likely causing the simple approximations (2.72) and (2.77) to fail. However, comparisons with the results of experiments and MC simulations have shown that the angular distribution predicted by the diffusion approximation has a surprisingly large range of validity. The equations (2.121) are particularly relevant for practical applications, since they state that both time-resolved and steady-state measurements of outgoing radiation are affected by the acceptance angle of the detection system by a factor that is independent of both time t and the exit time ρ . If needed, this factor can be calculated integrating $F(\cos \theta_e)$ on the numerical aperture of the detection system. However, for many applications the knowledge of this factor is not necessary, since measurements are available only in relative units. As long as the angular dependence of the outgoing radiance can be represented by equation (2.122), the solutions for the total outgoing flux can be used to analyze experimental data without introducing any approximation.

2.2.5 Validity of the diffusive approximation

The study of how accurate the diffusion approximation (DA) is and of its validity range has been going on since its origin and still represents an active research field, largely

motivated by the appealing power and simplicity of the diffusion theoretical framework [67–72]. The fundamental assumptions at the base of the DA are the validity of the P_1 expansion (2.72) and the assumption of stationary flux over time scales comparable to l'_s/ν (2.75). These assumptions are expected to fail very close to surfaces and sources, both of which would, in principle, invalidate the simple angular dependence imposed on $I(\mathbf{r}, t, \mathbf{s})$ and possibly the magnitude of $\partial F/\partial t$.

Nonetheless, many of these limits have proven to be quite flexible, and as a matter of fact, the breakdown of the DA can occur in different ways and to different degrees depending on the experimental observable of choice. Moreover, in the past decades, a continuous effort has been devoted to deriving and refining the theoretical framework of diffusion theory, in order to make it more robust, clarify the nature of its underlying assumption and extend its validity range [61, 63, 70, 73–77]. Indeed, when carefully applied, diffusion theory often yields very good experimental results, in most cases much better than what was expected. The only rigorously defined limit for the validity of diffusion theory, as was the case with the RTE, is that it does not make much sense to apply it to distances smaller of the average distance between scatterers (and its corresponding time scale). Additionally, as for the RTE, significant deviations will arise in cases where there is a high coherent contribution (which, in a scattering system, occurs typically near the source). For this reason, reflection measurements are more prone to deviate from the diffusion prediction since reflected light will contain a significant contribution coming regions/times closer to the source modulation, with features depending heavily on the adopted source model. In all our derivations, we assumed an isotropic point source placed at a depth $z_{\text{src}} = l'_s$ as an approximation of a pencil beam source coming from outside the material, which in turn is an approximation of an experimental beam. All these models become eventually indistinguishable both in transmission and reflection, provided that the measurement is performed sufficiently far and at a sufficiently large delay from the source emission [78, 79].

A separate but directionality-related problem is that of the validity of the similarity relation, which accounts to equating arbitrarily complex phase functions to a simpler one with the same first moment. When the anisotropy factor g approaches its upper limit (as it is the case for most biological samples, where often $g \simeq 1$ [6, 80]), the length scales over which the P_1 expansion can be considered valid grow proportionally casting a stronger constraint on the validity range of the DA, and can give rise to appreciable deviations especially near the boundaries [76, 81].

The interplay between the presence of absorption, the definition of the diffusion constant and the validity range of the diffusion approximation has been the subject of a long debate [57–59, 82], and as such it deserves a few words within the scope of this work. When looking at transport problems in the time domain, as it is the case for this thesis, the role of absorption becomes much clearer and *factorizes out of the problem*. This holds *irrespective of the validity of the diffusion approximation*, as long as coherent effects can be disregarded in the scattering system of interest. Regarding the diffusion approximation, its validity remains unaffected provided that we consider time scales that are orders of magnitude longer than l'_s/ν . As long that we have access to this time range (either experimentally or numerically), the presence of absorption does not affect

the validity of the diffusion approximation nor the definition of the diffusion constant. The situation is completely different in a steady-state configuration, where the presence of absorption directly causes low-order scattered light to dominate the integrated signal. This of course breaks the validity assumptions of diffusion, and the definition of the diffusion constant must be modified to take into account the presence of absorption [59, 74, 83]. In the general time-domain picture, however, diffusion and absorption can be considered as two unrelated parameters, as discussed in subsection 2.2.2. The growth of the spatial variance of the profiles, which lies at the very definition of the concept of diffusion, is completely unaffected by the presence of absorption, which can be therefore regarded as an irrelevant parameter as long as the validity of the diffusion approximation is concerned.

Based on these remarks, when considering the validity of the diffusive approximation in a slab geometry for transmitted light, an important figure of merit driving the transition from the diffusive to the ballistic regime is, unsurprisingly, the distance between the source and the distal boundary of the slab where the measurement is performed [70, 72, 84–86]. One common way to parametrize it is to introduce the so-called reduced optical thickness or simply optical thickness (OT) L/l'_s . The usual assumption is that the thickness of the medium should be at least one order of magnitude larger than the reduced transport mean free path (few authors give a rule-of-thumb factor of 8 [72, 87]). As a matter of fact, already at OT = 10 (to be compared with the illustrative example with OT = 100 shown in the previous subsection), small deviations from the DA start to appear as shown in Figure 2.10, where the simulated distributions for a sample with $L = 1$ mm, $g = 0.6034$ and $l_s = 39.66 \mu\text{m}$ ($l'_s = 100 \mu\text{m}$) are compared with the results of the DA equations. As can be seen, all observables exhibit deviations of different entity at early times, and qualitatively analogue results are obtained when looking at the reflectance. At high refractive index contrast, some deviations can persist also at late times, signaling the failure of the diffusive approximation. This happens in particular for the total transmittance $T(t)$, where diffusion theory is not able to reproduce correctly the asymptotic decay constant, while notably the mean square width growth is less affected in spite of the very slow convergence of the excess kurtosis towards 0.

Finally, the bottom panels of Figure 2.10 illustrate how the angular distribution of transmitted light evolves in time and eventually converges towards the distribution predicted by diffusion theory (2.122). In general, the distribution will exhibit a more or less pronounced forward peak towards $\cos \theta = 1$ due to the fact that early transmitted walkers/light will necessarily have gone through the sample perpendicularly. Few, well documented deviations at grazing exit angles remain even at long times for the index-matched case, due to the presence of a scattering anisotropy $g \neq 0$. The interesting thing to notice here is, however, that the final distribution becomes quickly stable, ceasing to evolve in time (and space), which is fundamental to guarantee that radiometric measurement taken on a finite aperture/detector are nonetheless representative of the intensity integrated over the whole solid angle.

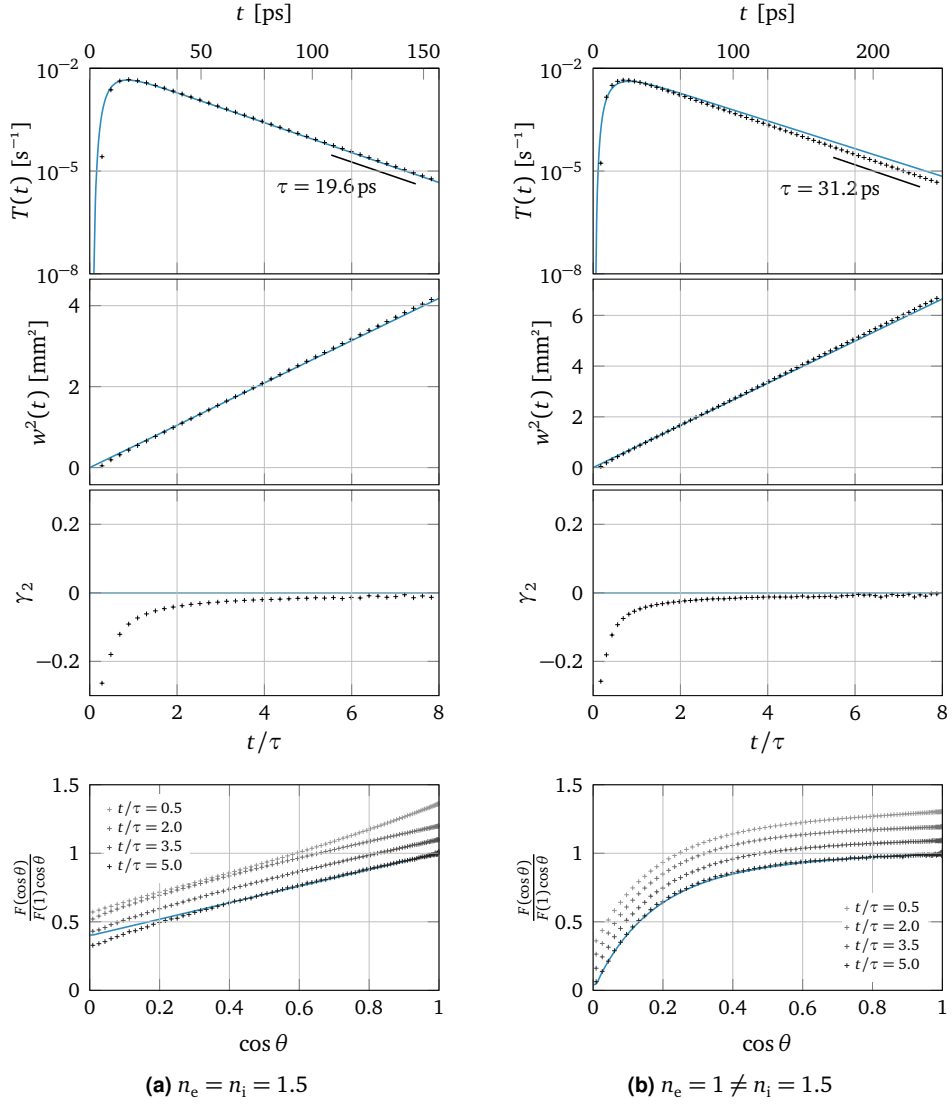


Fig. 2.10.: Breakdown of the diffusion approximation for a slab sample with reduced optical thickness $L/l'_s = 10$ for the (a) index-matched and (b) index-mismatched case. For the sake of brevity, the distributions obtained via MC simulations are compared with those predicted using the DA only for transmitted light. From top to bottom, the time-resolved flux at long times is poorly modeled with increasing refractive index contrast; on the contrary, the MSW is almost unaffected, even if the excess kurtosis vanishes slowly. The angular distributions converge quickly and accurately towards the DA prediction (simulated distributions have been shifted vertically for better visibility). As expected, early times are generally associated with higher deviations.

Spatially-resolved time-of-flight spectroscopy

3

” *If your experiment needs statistics, you ought to have done a better experiment.*

— E. Rutherford
(physicist)

3.1 Optical characterization of turbid media

3.1.1 Overview of experimental techniques

The quest to achieve accurate optical characterization of turbid media has driven the development of an array of experimental techniques in the past decades. In most cases the techniques are meant to be non-invasive and aim at retrieving absorption and scattering properties from measurements taken at the boundary of a target medium. These techniques can be either *direct* or *indirect*, depending on whether a particular microscopic coefficient is measured in a way that does not require a model of light propagation, or are obtained by solving an inverse problem. In the latter case, the optical properties are plugged into a suitable light propagation model as fitting parameters.

Photometric techniques can be classified in terms of the quantity being measured (intensity, fluence rate or radiance) and, most importantly, on the domain in which the measurement is resolved. Several examples exist ranging from steady-state integration, to measurements in the spatial, angular, temporal and frequency domain. Each of these cases is shown schematically in Figure 3.1, qualitatively ranked from left to right in order of increasing information content according to the technical limitations of each technique [88, 89].

The first three cases belong to the family of continuous wave (CW) (or intensity) techniques, which are usually associated with fluence-type measurements. In principle, the simplest configuration is that of using a single or two joint integrating spheres to measure the total reflected, transmitted and collimated light (which can be let out of the sphere through a small aperture) [90, 91]. Albeit simple, integrating spheres are prone to systematic errors, and the need to acquire absolute fluence levels requires cumbersome and impractical calibration of the setup [92–94]. Radiance techniques, where a detector with well-defined angular aperture is rotated around the sample to record the angular distribution of re-emitted light, improve on the integrating sphere technique by adding further information to recover the optical properties, especially as




Fig. 3.1.: From left to right: classification of available experimental techniques for optical characterization of turbid media in terms of the measurement domain. An external stimulus is shown on the left side of a block representing the investigated medium. The initial (typically narrow) distribution undergoes a modification in some domain following interaction with the sample, resulting in an attenuated, broader signal on the distal side of the block.

regards the scattering anisotropy [95, 96]. In comparison, spatially resolved steady-state is much more popular among CW techniques due to its simplicity and low cost [78, 97]. The spatial decay of fluence from the source, though, depends both on the absorption and scattering coefficients and is not separable unless, again, absolute fluence levels are recorded. Alternatively, measurements should be taken close to the source, where the diffusion approximation breaks down. Other steady-state schemes involve patterned illumination, which can conveniently provide single-shot, wide-field characterization of the local properties of scattering as seen from outside a medium [98, 99]. In general, however, the common downside of CW techniques is that their accuracy and validity is usually restricted to a limited range of optical parameters and sample geometries.

The second main class of techniques exploits intensity modulation of the source. In the frequency domain, this is done by recording the diffuse waves propagating when the source is modulated at a typical frequency between 100 MHz and 1 GHz [100, 101]. In response to the external modulation stimulus, the medium attenuates the signal and introduces a phase shift relative to the input signal. On the other hand, in time resolved techniques, a short pulse of light is launched into the medium at a selected point and time. As it propagates through the medium, the pulse undergoes attenuation and broadening that can be used to extract parameters through fitting of a suitable forward model [56, 102]. It should be noted that in principle the time and frequency domains are formally equivalent and simply related to each other via a Fourier transform operation. However, different practical limitations exist associated to both techniques, determining their preferred use depending on the situation. Frequency-domain measurements typically offer lower noise compared to time-resolved measurement, but on the other hand the information content is lower unless the measurements are performed at many modulation frequencies. The maximum bandwidth is also usually lower for a frequency-resolved system than a time-resolved, which is important especially for small distances or sample geometries. In this respect, time-domain approaches are most comprehensive, since a short laser pulse in the fs range implicitly contains all the modulation frequencies, including the zero-frequency component [103]. In comparison, modulation frequencies of the order of a GHz correspond to a resolution of the order of a ns and to density wavelengths of the order of tens of cm and therefore are unsuited to study smaller specimens or the finer details of transport.

The major advantage of time-resolved measurements is that the analysis of the pulse shape, rather than the fluence levels, obviates the need for absolute calibration. Even though the equipment needed to perform a time-domain experiment is more complex and costly than CW-based methods, time-of-flight measurements allows straightforward separation of scattering and absorption contributions from a single measurement, and are virtually capable of handling any combination of optical properties and sample geometry, provided that a suitable model of light propagation is used.

In this chapter we will describe and review the components and performance of a time-domain characterization technique and its major upgrade to include spatially-resolved information. As a matter of fact, in recent years, a clear trend is emerging to take advantage of multi-domain optical characterization, even if partially, and the rigid classification of experimental techniques in terms of their domain is becoming fuzzier. Notable examples include mixing of spatial and angular domains (e.g., exploiting spatial filtering, collimated detection or more complex configurations [89, 104]) or of spatial and temporal information, which has been typically done by moving a collection fiber at different points of a sample [105, 106]. In most cases, though, only few bits of cross-domain information are collected, or different single-domain measurements are compared separately. As we will demonstrate through this and the next Chapter, gaining access to the full spatial and temporal information enables new strategies for accurate optical characterization that are not possible with partial multi-domain information or multiple single-domain measurements.

3.1.2 Time-domain techniques and sub-ps optical gating

Several experimental configurations are documented in the literature to implement time-of-flight measurements, each with different advantages and typical time scales of operation. The most common technique is represented by time-correlated single-photon counting (TCSPC), which has the advantage of working at very low intensity levels and is based on individual photon statistics. A constant fraction discriminator is used to accurately determine the arrival of pulses with a resolution that can be around few tens of picoseconds, well below the rise time of the detector [56, 107, 108]. Gated CCD cameras have also been used to obtain time-resolved detection systems, having the major advantage of yielding parallel measurements over a wide field of view, even though with a temporal resolution limited to few hundreds of ps [109, 110]. On the other hand, a resolution of the order of just few ps is typically obtainable using streak cameras, which have been also largely used in the past for optical characterization of turbid media [69, 111, 112], despite their high cost and limited dynamic range.

A completely different approach is needed to study transport on much shorter time scales, a challenging experimental task which rules out electronic-based devices because of their inherently limited response time. Several ultrafast shuttering techniques have been exploited in the past decades to effectively time-gate reflected or transmitted light at fixed delays, either by amplifying it or attenuating the signal at different exit times. From a historical perspective, the development of these techniques was mainly driven by transillumination applications, which have the sole aim of *suppressing* scattering

from turbid media in order to see through them, rather than to characterize it. For this reason, a large part of the literature available on ultrafast gating techniques is influenced or is somehow related to this separate field [113, 114]. Quite surprisingly, the vast characterization potential of sub-ps gating wide-field techniques was never exploited, to our knowledge, to obtain a multi-domain characterization of the optical properties of scattering media, supposedly because of the difficulties in controlling or even assessing their quantitative reliability.

Several strategies exist in order to achieve an all-optical gating operation. A first notable example is that enabled by the Kerr effect. A Kerr cell placed in between crossed polarizers can be in fact operated as a fast shutter triggering its birefringence by means of short laser pulses [115, 116]. A Kerr gate can be thought of as the analogue of a mechanical shutter in the ps range. As such, it is not wavelength nor angle dependent, and in fact it can be used to collect light with broad spectral and angular distributions. Unfortunately, its performance is ultimately limited by the dynamic range of the transmission opacity of the cell, which can hardly exceed 10^4 , meaning that only sufficiently intense signals can be measured. Similar limitations exist for amplification obtained by means of stimulated Raman scattering (SRS) occurring in materials such as hydrogen gas, where a long-wavelength (Stokes) beam is amplified by a shorter-wavelength pump beam [117].

Turning to high-gain ultrafast amplifying gates, applications are mostly based on second-harmonic generation (SHG) or, more in general, on optical parametric amplification (OPA) [117–123]. Different laser pulses can be combined spatially and temporally on certain nonlinear crystals to generate light at frequencies equal to the sum and difference of the incoming beams. Typically, the frequency upconversion occurring in sum-frequency generation (SFG) can be profitably exploited to perform the experiment at longer wavelengths of interest (typically in the near infrared range) while detecting the resulting signal at shorter wavelength at which detectors are more efficient. Despite the high gain and excellent temporal resolution, however, these optical gating techniques have been often overlooked by the transillumination community because of the extremely narrow angular selection dictated by the phase-matching condition. As we will discuss in the following sections, this rigid constrain does not pose a significant limitation when studying light transport in turbid media.

3.2 Cross-correlation optical gating

In this Section we describe the experimental properties and working principles of an optical apparatus based on the cross-correlation gating technique [118, 119]. In a typical optical gating apparatus, represented schematically in Figure 3.2, two synchronous, collinear probe and gate pulses are made to impinge on a nonlinear crystal. The probe pulse at frequency ω_1 (interpreted as the central frequency of the pulse bandwidth) impinges on the scattering sample and interacts with its structure, emerging with a broadened intensity distribution $I_{\omega_1}(t)$. Conversely, the gate pulse propagates unaltered in free space preserving its original temporal profile $I_{\omega_2}(t)$. When the two pulses eventually reach the crystal, an upconverted signal at frequency $\omega_1 + \omega_2$ is generated




Fig. 3.2.: Working principle of an optical gating setup. A sub-ps probe pulse (in red) with central frequency ω_1 is broadened in time following interaction with a scattering medium. The resulting broadened signal impinges on a nonlinear crystal (NLC) together with a narrow gate pulse (grey). The relative delay between the two can be tuned using a translation stage. At each position/delay, a sum-frequency signal is generated at the crystal with an intensity proportional to their convolution integral (3.1).

depending on the degree of spatial and temporal overlap. This frequency mixing process is known as sum-frequency generation (SFG) or upconversion and occurs in nonlinear crystals with finite second-order susceptibility $\chi^{(2)}$. The intensity profile of the sum-frequency signal is given by the convolution integral of the intensities

$$I_{\omega_3}(\Delta\tau) = \int_0^{\infty} I_{\omega_1}(t) I_{\omega_2}(t - \Delta\tau) dt \quad (3.1)$$

where $\Delta\tau$ represents the relative delay between the two pulses, which can be set using a translation stage. For a fixed delay, the resulting train of signal pulses $I_{\omega_3}(\Delta\tau)$ has now a stationary intensity which can be detected and integrated with a slow detector. By scanning over the delay $\Delta\tau$ it is possible to sample the temporal evolution of $I_{\omega_1}(t)$ using the unperturbed $I_{\omega_2}(t)$ pulse as a temporal gate (hence the name of the technique), by performing a deconvolution operation. The problem simplifies if we can assume that $I_2(t)$ has a much shorter duration than any signal it is convoluted with, as it is well the case for all the measurements presented in this work. In this case, the gate pulse can be approximated by a $\delta(t)$ pulse, yielding

$$I_{\omega_3}(\Delta\tau) \simeq I_{\omega_1}(\Delta\tau) I_{\omega_2}(0) \quad (3.2)$$

meaning that the time evolution of the sum-frequency signal represents directly that of the investigated signal. On a side note, equation (3.2) also shows an interesting point: the intensity of the measured signal I_{ω_3} (and hence the signal-to-noise ratio) can be enhanced by transferring energy from either the probe or the reference pulse (which does not interact with the sample).

The temporal resolution that can be obtained with this setup is influenced by the width of the gate pulse (if its temporal shape is unknown and therefore cannot be deconvolved) and on the minimum spatial displacement that can be accurately performed by the translation stage. Displacements of $1\mu\text{m}$ can be easily obtained (corresponding to

an increase of the total path length of $2\mu\text{m}$), yielding a sampling rate of $\sim 6.7\text{ fs} \sim 1.5 \times 10^{14}$ samples/s. Moreover, by using a probe and gate pulse pair generated synchronously by the same laser source, the final resolution is exactly unaffected by fluctuations in the repetition rate or timing jitter of the pulses, which makes the setup inherently more stable and robust to undesired drifts.

3.2.1 Phase-matching and angular acceptance

The upconversion process at the core of optical gating is a delicate process which depends critically on the phase-matching between the interacting pulses, as required by the conservation of momentum

$$\Delta \mathbf{k} = 0 \quad \rightarrow \quad \mathbf{k}_{\omega_1} + \mathbf{k}_{\omega_2} = \mathbf{k}_{\omega_3}, \quad (3.3)$$

where the \mathbf{k}_{ω_i} are the wave vectors of the mixing beams at angular frequency ω_i . Equation (3.3) becomes scalar in the simple case of collinear geometry, yielding

$$\omega_1 n(\omega_1) + \omega_2 n(\omega_2) - (\omega_1 + \omega_2) n(\omega_3) = 0, \quad (3.4)$$

where $n(\omega_i)$ is the refractive index of the crystal at ω_i . It is generally not possible to satisfy the condition (3.4) in isotropic media or in centrosymmetric solids, but one can exploit the birefringence of some anisotropic crystals to obtain the right combination of refractive indexes depending on the angle of incidence and polarization of the incoming beams. In a uniaxial birefringent crystal with optic axis \hat{z} , both ordinary (polarization in the $\hat{x}\hat{y}$ plane) and extraordinary (polarization perpendicular to the $\hat{x}\hat{y}$ plane) rays are defined, seeing refractive indexes of n_o and n_e respectively. Each intermediate combination will yield a correspondingly weighted value of the effective refractive index along that direction, which allows to satisfy (3.4) by tuning the angle of incidence. Uniaxial crystals are classified as ‘negative’ if $n_e < n_o$ and ‘positive’ otherwise. In the former case the possible polarization combinations are $o + o \rightarrow e$ or $o + e \rightarrow e$ with the sum-frequency necessarily polarized along the extraordinary direction, while for positive crystals it must be ordinarily polarized, i.e., $e + o \rightarrow o$ or $e + e \rightarrow o$. Irrespective of the sign of the birefringence, the interaction is termed of type I if the incident beams have parallel polarization, and type II otherwise.

In our experimental setup, we use a square $5\text{ mm} \times 5\text{ mm} \times 2\text{ mm}$ β -Barium borate (BBO) crystal (Figure 3.3a). BBO is a negative crystal, which we use in a $o + o \rightarrow e$ (type I) configuration. The angle $\hat{\theta}$ between the optic axis and the incidence direction must therefore satisfy the phase-matching relation

$$\omega_1 n_o(\omega_1) + \omega_2 n_o(\omega_2) - (\omega_1 + \omega_2) n(\omega_3, \hat{\theta}) = 0, \quad (3.5)$$

with

$$n(\omega_3, \theta) = \sqrt{\frac{n_o^2(\omega_3) n_e^2(\omega_3)}{n_o^2(\omega_3) \sin^2 \theta + n_e^2(\omega_3) \cos^2 \theta}}. \quad (3.6)$$

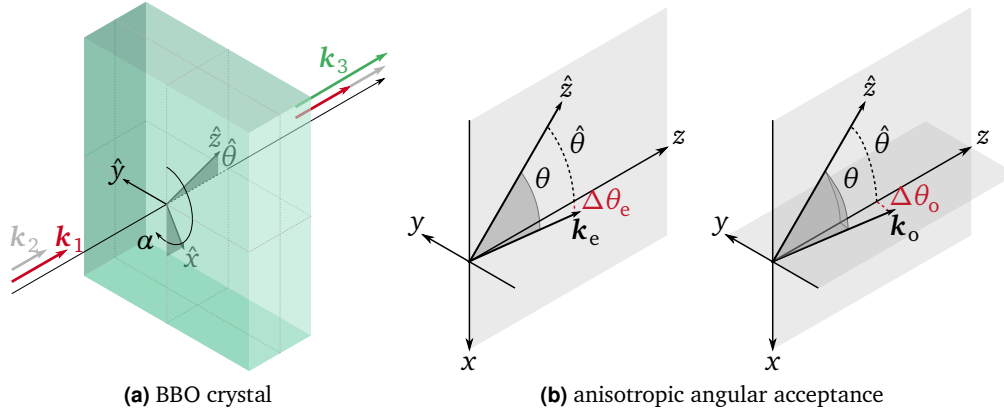


Fig. 3.3.: (a) Reference frame for a uniaxial birefringent crystal in a sum-frequency interaction (not to scale). The optic axis of the crystal lies along \hat{z} . (b) Geometric sketch showing how the phase-matching condition results in different angular acceptances along the ordinary and extraordinary directions. A slightly misaligned \mathbf{k}_e and a largely misaligned \mathbf{k}_o may result in the same angular mismatch $\Delta\theta$.

For a BBO crystal the relevant refractive indexes can be calculated using empirical Sellmaier's equations with coefficients

$$n_o^2 = 2.7405 + \frac{0.0184}{\lambda^2 - 0.0179} - 0.0155\lambda^2$$

$$n_e^2 = 2.3730 + \frac{0.0128}{\lambda^2 - 0.0156} - 0.0044\lambda^2$$

which give the required $n(\omega_3, \hat{\theta})$ through equation (3.5) for the specific pair of probe and gate frequencies used, which, in our experimental case, correspond to $\lambda_1 = 810$ nm and $\lambda_2 = 1.51$ μ m ($\lambda_3 = 527$ nm). Our crystal is cut at an angle $\hat{\theta} = 22.96^\circ$ in order to give maximum conversion efficiency for this combination of wavelengths at perpendicular incidence. If different wavelengths are used, the incidence angle can be adjusted accordingly by slightly tilting the crystal appropriately.

Once that the correct geometry is set, a certain efficiency is associated to the upconversion process, which can be defined through the generated power [124]

$$P_{\omega_3} = \eta P_{\omega_1} P_{\omega_2} \quad (3.7)$$

and depends critically on the phase mismatch $\Delta k = |\Delta \mathbf{k}|$, with a typical decay behavior

$$\eta(\Delta k) = \eta(0) \frac{\sin^2(L\Delta k)}{(L\Delta k)^2}, \quad (3.8)$$

where L represents the thickness of the nonlinear crystal. The efficiency will drop by a factor 2 when $L\Delta k \simeq \pm 1.3916$ rad, which helps us defining a spectral and an angular acceptance bandwidths beyond which the upconversion efficiency falls off. The frequency bandwidth is due to chromatic dispersion, and is related to the group velocity mismatch

of the mixing waves which will reduce the spatial overlap between the two pulses as they propagate along the crystal. If we consider, for example, a variation of the probe frequency, the associated full bandwidth can be calculated as (ω_3 varies automatically with ω_1 to preserve energy conservation)

$$\frac{\partial \Delta k}{\partial \omega_1} = \frac{\partial k_{\omega_1}}{\partial \omega_1} - \frac{\partial k_{\omega_3}}{\partial \omega_3} = \frac{1}{v_{g,1}} - \frac{1}{v_{g,3}} \rightarrow \Delta \omega_1 = \frac{2.7831}{|v_{g,3}^{-1} - v_{g,1}^{-1}| L}, \quad (3.9)$$

which, comparing the group indexes at our probe and gate frequencies, yields a relative delay of roughly 90 fs along a crystal length of 2 mm, corresponding to a phase-matching bandwidth of ~ 10 nm [125]. For this reason, a 2 mm-thick crystal represents a sensible choice for pump and gate pulses of roughly the same bandwidth and duration (compare Figure 3.4). Similar considerations hold for the angular acceptance, with a few differences. In a type I configuration, the spatial walk-off related to the angular acceptance affects only the (extraordinary) sum-frequency signal, leaving the total interaction length between the probe and gate pulses unmodified. In turn, this slightly reduces the beam quality of the sum-frequency beam, which will be slightly elongated along the walk-off direction and therefore affect the spatial point spread function. The walk-off angle ϱ for the sum-frequency beam can be calculated using the definition of $n(\omega_3, \theta)$ (3.6) as

$$\varrho = -\frac{1}{n(\omega_3)} \frac{\partial n(\omega_3)}{\partial \theta} = \frac{(n_e^2(\omega_3) - n_o^2(\omega_3)) \sin \theta \cos \theta}{n_o^2(\omega_3) \sin^2 \theta + n_e^2(\omega_3) \cos^2 \theta} \quad (3.10)$$

which yields $\varrho \simeq 3^\circ$, corresponding to a broadening of the sum-frequency spot at the exit surface of the 2 mm crystal to a width of roughly 100 μm . Finally, an important role is played by the phase mismatch induced by an angular misalignment between the probe and the gate pulse. This is of particular relevance for our application since, while our gate beam propagates collimated to the crystal, our probe signal will comprise multiply scattered light distributed over a broad range of incoming angles (see, for example, the simulated distributions of Figure 2.10). This means that we will effectively upconvert only a fraction of the light transmitted by our samples, depending on the angular filtering actuated by the phase-matching condition. The full-angle-half-maximum width can be calculated similarly to the frequency case, considering a mismatch θ between the probe and the gate beams

$$\frac{\partial \Delta k}{\partial \theta} = \frac{\omega_3}{c} \frac{\partial n(\omega_3, \theta)}{\partial \theta} \rightarrow \frac{\partial \Delta k}{\partial \theta} = -\frac{\omega_3}{c} n^3(\omega_3, \theta) \left[\frac{1}{n_e^2(\omega_3)} - \frac{1}{n_o^2(\omega_3)} \right] \sin \theta \cos \theta \quad (3.11)$$

which leads to

$$\Delta \theta = \frac{2.783c}{L \omega_3 |n^3(\omega_3, \theta) [n_e^{-2}(\omega_3) - n_o^{-2}(\omega_3)] \sin \theta \cos \theta|} \simeq 0.1^\circ. \quad (3.12)$$

As can be seen, the value of the angular acceptance for our experimental configuration is extremely small using our experimental wavelengths. This is a well known property of

critical phase-matching optical gating, which in fact found application as a transillumination technique where the narrow collinear filtering can be beneficial to the selection of ballistic light.

From an experimental point of view, it is important to be aware of the possible effects of this spatial filtering when comparing measured data with a given forward model. Two points in particular are worth discussing. The first one regards the fact that the value that we obtained for the full-angle-half-maximum $\Delta\theta$ refers naturally to the $\hat{x}\hat{y}\hat{z}$ crystalline reference frame, which is tilted by an angle $\hat{\theta}$ with respect to the faces of the crystal. This means that a certain misalignment in the xyz reference of the plane surfaces of the crystal might result in very different values of $\Delta\theta$ depending on whether we are moving away from $\hat{\theta}$ tangentially to it. This difference can be qualitatively appreciated by looking at the sketches shown in Figure 3.3b, and will introduce an anisotropic angular acceptance in the xy plane of the crystal. This effect must be taken into account and corrected for when moving to upconversion imaging applications, as we will discuss further in the following subsections.

The second point concerns how this narrow angular selection affects evaluation of time-resolved data depending on the model used. In the literature, when the diffusive approximation is used to evaluate a spatially-integrated time-resolved dataset obtained with a similar setup, it is always tacitly assumed that there is a simple proportionality between the total time-resolved transmittance and any angularly filtered fraction, as shown explicitly in subsection 2.2.4. However, at early times this assumption fails more markedly in proximity of the direction of collimated transmission, which is the only one that is effectively upconverted. While this is not a problem when the diffusive approximation is used (early transmitted light should be discarded anyways), care has to be taken also when using the gold-standard method of Monte Carlo simulations. As a matter of fact, it has been demonstrated that the breakdown of the diffusive approximation occurring at early times when measuring in transmission can be profitably exploited to retrieve the otherwise degenerate scattering anisotropy factor g [102]. However, as it is apparent from the simulations shown in Figure 2.10 for a sample of comparable thickness, the shape of the angular distribution at early times is fairly time-dependent, which should be taken in account also in the forward model, e.g., by considering only the walkers that are transmitted within a narrow angular range close to $\cos\theta = 1$. As regards the results presented in this thesis work, experimental data are compared with MC simulations performed with a custom software suite (described in Chapter 5) capable of taking into account arbitrary angular filtering conditions. As a matter of fact, however, all the data and numerical evaluation techniques that we describe in the following are based on the late-time behavior of transmittance, where we have tested that the net effect of the angular selection applied by the effective numerical aperture of our optical gating setup is completely negligible.

3.2.2 Sources characterization

A detailed characterization of the laser sources used in the experiment is of paramount importance to assess the correct functioning of the gating application, its temporal

resolution, and the correct modeling of the source term in the direct and inverse models of light transport discussed in the previous chapter. In our experiment, the probe and gate pulses are at two different wavelengths. Although the working principle of optical gating is still valid in a second-harmonic generation scheme, employing different wavelengths offers several practical advantages such as the compatibility with a collinear illumination scheme and the ability to easily remove background signal spectrally. Moreover, the ‘probe’ and ‘gate’ roles are perfectly interchangeable between the two arms, that are perfectly symmetric. This allows to investigate any sample in two different wavelength regions, enabling a multi-spectral characterization without the need to modify the setup.

The laser sources employed in our experiment are arranged as follows: a Spectra-Physics® *Millennia* diode-pumped solid-state (DPSS) CW laser emitting at 532 nm with a power of 8.6 W pumps a *Tsunami* Ti:Sa mode-locked laser which generates a train of pulses at $\lambda_1 = 810$ nm with a repetition rate of 82 MHz and an average power of 1.65 W. The output of the *Tsunami* is fed into an *Opal* optical parametric oscillator (OPO) to yield a pair of downconverted beams of which only the shorter wavelength one ($\lambda_2 = 1.51 \mu\text{m}$) is used as the gate beam (see Figure 3.5). Different frequencies can be obtained if desired via temperature tuning of the downconverting crystal, in the 1.4–1.6 μm range. The probe beam, conversely, is provided by the non converted residual at 810 nm. The average power of the two beams at the output apertures of the *Opal* is of 0.20 W and 0.28 W, respectively.

A temporal characterization of the probe pulse and a spectral characterization of both beams is performed at the beginning of each measurement session using a compact autocorrelator based on a Mach-Zehnder interferometer [126] and a spectrum analyzer. The expected pulse shape for an actively mode-locked Ti:Sa laser is given by a hyperbolic-secant function

$$f(t; \Delta\tau_p) = \sqrt{\frac{\gamma_p}{2}} \text{sech}(\gamma_p t) \rightarrow I(t; \Delta\tau_p) = \frac{\gamma_p}{2} \text{sech}^2(\gamma_p t), \quad \gamma_p = 2 \ln(1 + \sqrt{2})/\Delta\tau_p \quad (3.13)$$

where $\Delta\tau_p$ is the full-width-half-maximum pulse duration. The sech^2 pulse shape arises, in spite of the active mode-locking, because of the pulse compression provided by the negative group delay dispersion added in the cavity to make it stable [127, @128]. The autocorrelation of two identical sech^2 pulses can be derived analytically and is given by

$$I_{ac}(t) = \gamma_{ac} \text{csch}^2(\gamma_{ac} t) [\gamma_{ac} \coth^2(\gamma_{ac} t) - 1], \quad \gamma_{ac} = 2 \ln(1 + \sqrt{2})/\Delta\tau_{ac} \quad (3.14)$$

which we can use to fit the envelope of the autocorrelation trace (Figure 3.4a). The full-width-half-maximum of the convoluted signal is equal to 144.6 fs, corresponding to a probe pulse width of $\Delta\tau_p = 93.7$ fs. The pulse duration can be compared with its spectrum, which also exhibits a sech^2 dependence as can be found by a Fourier transform operation

$$\tilde{I}(\omega; \Delta\tau_p) = \frac{\pi^2}{2\gamma_p} \text{sech}^2 \left[\frac{\pi}{2\gamma_p} (\omega - \omega_i) \right]. \quad (3.15)$$

Figures 3.4b and 3.4c show typical spectra measured for the probe and gate beams, and their respective fits with equation (3.15), yielding a central wavelength of 810.2 nm and

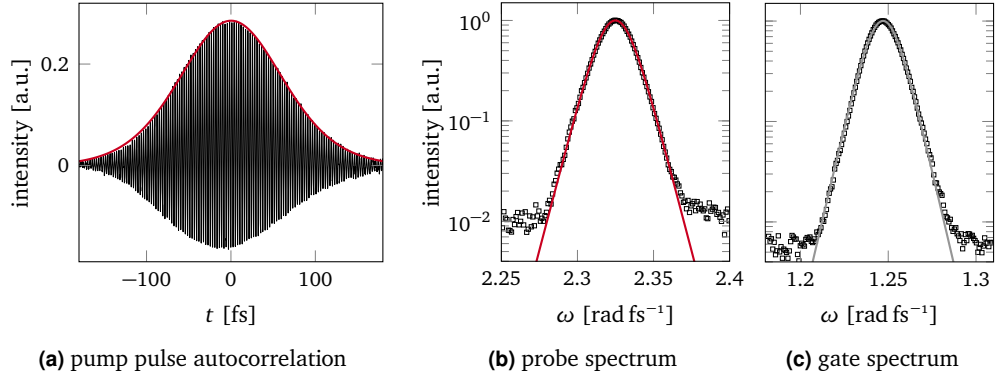


Fig. 3.4.: (a) Autocorrelation measurement of the residual pulse. The envelope profile fringes are fitted with the autoconvolution of two $\text{sech}^2(t)$ pulses. (b) Spectrum of the unconverted residual (probe pulse) from the OPO, with a $\text{sech}^2(\omega)$ fit. (c) Spectrum of the OPO signal (gate pulse) with a $\text{sech}^2(\omega)$ fit.

1510.5 nm, respectively. The values of $\Delta\tau_p$ obtained from fitting the spectra correspond to their associated Fourier limited pulse widths, which, in the case of the probe, would predict a full-width-half-maximum duration of 74.7 fs. Comparing this value with that obtained by the fit of the autocorrelation trace reveals that our experimental pulses are not strictly Fourier limited. Regarding the gate pulse, its duration cannot be measured directly, but we can nonetheless retrieve it by adapting iteratively the a cross-correlation time trace to the numerical convolution of the (known) pump pulse with a sech^2 pulse of unknown duration. The resulting fit is shown in Figure 3.6a, returning a gate pulse duration of 134 fs, to be compared with a typical full-width-half-maximum of the whole cross-correlation signal of roughly 170 fs. This value can be thought of as the analogue of the instrument response time, and we assumed it as the duration of a pulsed sech^2 excitation in all our Monte Carlo simulations to be compared to experimental data.

3.3 Ultrafast imaging

The experimental setup that we developed to perform transmittance and reflectance measurements both in the temporal and spatial domain builds upon a preexisting optical gating setup described in [129], which has been improved and modified for this purpose as shown in Figure 3.5. The setup is shown in a transmission configuration, but reflectance can be measured as well by rearranging the illumination and collection geometry accordingly.

Several examples exist in the literature where the cross-correlation technique has been used to characterize ultrafast pulse propagation [130, 131], generally disregarding the spatial distribution of the converted signal (Figure 3.6a). Nevertheless, cross-correlation gating offers several advantages, such as being unaffected by optical chirp (in contrast with interferometric methods) and being compatible with a collinear excitation scheme because spurious second harmonic contributions from gate/probe beams, if any, can still be spectrally separated. Moreover, a collinear geometry is particularly convenient to




Fig. 3.5.: Schematic of the experimental setup. The output of a mode-locked Ti:Sa laser is fed into an optical oscillator. The unconverted residual (red) and the synchronous downconverted signal (grey) are used as the probe and gate pulses, respectively. The gate beam is collimated and expanded to guarantee uniform illumination of the BBO crystal. A double-telecentric system (L_2 and L_3) collects the light scattered from the sample, a diaphragm (FF) imposes isotropic angular acceptance. The gate and probe beams are superimposed with a dichroic mirror (DM). The crystal and the detectors are enclosed into a shielded box. A long-pass filter (LPF) at the entrance of the box ensures that no visible light can enter from outside. Time-resolved detection can be performed either with a photo-multiplier tube (PMT) or with a CCD camera. A chopper on the probe arm enables automatic background subtraction when using a gated photon counter.

mitigate possible distortion of upconverted images passing through the BBO, which is a crucial aspect for accurate spatially resolved experiments [132]. Still, in order to adapt the experimental configuration to an ultrafast imaging application, two main modifications have been adopted to further improve its performance and quantitative accuracy, namely the addition of a double-telecentric optical apparatus and of a tunable Fourier-space filter (diaphragm FF). Together, they allow to transfer the paraxial component of transmitted light unaltered to the BBO crystal, and to correct for its anisotropic angular acceptance.

The double telecentric system is implemented using two identical lenses (L_2 and L_3) in a $4f$ configuration, with $f = 100$ mm. Different lenses can be used if a higher magnification is needed. In our case, the resulting magnification was approximately of $1\times$, allowing to cover a field-of-view of a few mm in size. With respect to a fixed focal lens, a double telecentric system offers lower distortion over the field of view, but most importantly preserves the angles between the object and the image. As we have seen in subsection 3.2.1, only an extremely narrow range of \mathbf{k} vectors falls within the angular acceptance determined by the phase-matching condition. In order to have a constant upconversion efficiency over the whole field-of-view, which is fundamental to achieve a high quantitative accuracy, the angular selection that is applied on the crystal (image) should be independent of the exit point on the surface of the sample (object). Having both the entrance and exit pupils at infinity guarantees that light that impinges perpendicularly on the BBO (and is therefore upconverted with maximum intensity) has left the sample with the same perpendicular angle irrespective of its exit point.

The presence of a tunable aperture in the back-focal plane of lens L_2 allows to vary the numerical aperture of the collection optics, which in turn affects the spatial resolution of

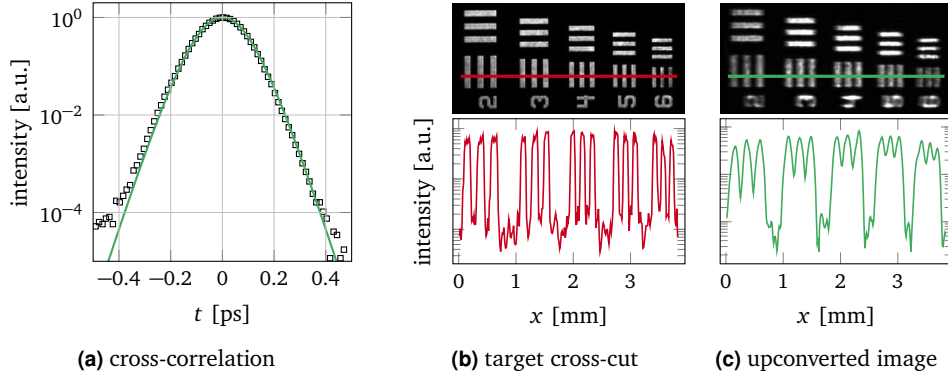


Fig. 3.6.: (a) Typical instrument response function obtained as the cross-correlation signal of the probe and gate pulses propagating with no intervening scattering sample nor lenses. A fit with the numerical convolution of two sech^2 pulses returns a full-width-half-maximum of 174 fs. (c) Switching the PMT detector with the CCD camera, the spatial information encoded in the signal is retrieved, showing the spatial resolution achieved in the upconversion process. The resolution of the imaging system under direct illumination is displayed in (b) for comparison.

the system. As we have seen, the phase-matching condition already determines a rigid constrain on the angular acceptance of the sum-frequency spatial frequencies, which is very narrow and anisotropic along the xy image plane. A few words should be spent on how this can alternatively impact the spatial resolution or the spatial uniformity of upconverted image, depending on the experimental configuration. Several examples are reported in the literature regarding full-frame imaging applications [120, 133–140], most of which employ the non-linear crystal in the Fourier plane of the imaging optics rather than in the image plane. Both configurations present different advantages and their adoption is to be preferred depending on the specific application. In the common Fourier-plane configuration the angular acceptance limits the upconversion uniformity over the whole field of view, with a decreasing efficiency with increasing distance from the optical axis, while in our case the image is formed with a cut-off set of k vectors (which lowers the resolution) while guaranteeing a spatially uniform up-conversion efficiency. Despite of its low quantitative fidelity, the Fourier configuration is typically preferred because of the sharper images that it delivers, and because it allows to use a more tightly focused gate beam which enhances the intensity of the sum-frequency signal. Conversely, we adopted an image-plane configuration since our diffuse signal typically does not exhibit any fine spatial feature, while a spatially uniform response represents a critical condition to attain our quantitative accuracy goal over multiple orders of magnitude.

The effect of the spatial frequency cut-off can be appreciated by comparing the images of a USAF 1951 resolution target obtained under direct illumination with a collimated probe beam (Figure 3.6b) and its transient, sum-frequency replica (Figure 3.6c). As can be seen, the overall resolution of the upconverted image is degraded due to the narrow filtering, and it is slightly worse along the x axis (compare with Figure 3.3b) where the cut-off sets in more steeply. The spatial resolution obtained along the two directions is

of 11.3 lp/mm along x and 22.6 lp/mm along y , calculated at a contrast threshold of 0.5. An isotropic angular acceptance of 11.3 lp/mm is finally achieved by closing the diaphragm (FF) to an aperture below 1 mm, which is anyway comparable to what is typically achieved in the literature [120, 136, 139, 140].

The rest of the setup is organized as follows: the laser sources, shown in the upper part of Figure 3.5, are characterized using a spectrum analyzer and an autocorrelator (not shown) as discussed previously. The unconverted residual from the Ti:Sa laser is used as the probe beam while the OPO signal serves as the gate (the idler is discarded). The gate and the probe pulses are perfectly interchangeable between the two arms of the setup, as well as the position of the sample and of the delay line. A focusing lens L_1 focuses the probe beam on the entrance surface of the sample. The size of the focal spot does not represent a particularly relevant parameter in the case of homogeneous turbid media. In fact, total time-resolved transmittance curves do not depend on the excitation spot position (and therefore on its size), and would yield the same result also for a plane wave excitation. On the other hand, when considering the mean square width expansion of the spatial profiles, a larger excitation spot would shift the retrieved values, but leave the expansion rate unmodified. Therefore, the only strict requirement regarding the size of the excitation is that the transmitted profile should fit within the field of view of the collection optics at any delay. For the sake of convenience and for simpler modeling, we decided to work with a rather tight focus of $\sim 10\ \mu\text{m}$, in order to assume a pencil beam configuration in the diffusive approximation. MC simulations are performed with a spatially gaussian profile of the actual experimental size. As regards the gate beam, a half-wave plate allows to turn its polarization to ensure that it matches that of the probe beam as required in a type I upconversion scheme. The beam is subsequently collimated and expanded to a size that is much larger than the crystal free surface, in order to obtain a uniform upconversion efficiency over the whole field-of-view. If properly collimated, the gate beam can be approximated as a plane wave impinging on the BBO with a perpendicular wave-vector.

The intensity of the sum-frequency generated signal typically spans a very large dynamic range, and the detection part of the setup must be shielded against ambient light to avoid spurious signal in the photon counting regime. Detection is performed either by using a photo-multiplier tube (PMT) in a spatially-integrated configuration (see, for example, Figure 3.6a) in combination with a gated photon counter and a chopper for automatic background subtraction, or by taking an image with a back-illuminated Andor iKon M912 CCD camera (Figure 3.6c). Inside the detection box, two independent imaging systems (not shown) can be switched to form an image of the BBO surface optimized for the probe (for alignment purposes) or the sum-frequency wavelength. Isolation of the sum-frequency wavelength is obtained with a long-pass filter (LPF) at the entrance of the detection box, and using a cascade of a dichroic mirror and band-pass filters obtaining an optical density exceeding 15 at the probe wavelength and above 10 at the second harmonic of the probe and third harmonic of the gate. The delay line and the sample positioning systems are motorized and controlled by a PC routine. Several sequences can be easily set to repeat automatically the measurement multiple times at

different positions of the sample, or to move it while integrating the signal at each fixed delay to obtain a measurement that is more representative of its average properties.

To conclude, it is interesting to comment on the overall sensitivity of the upconversion process in terms of absolute fluence. Previous measurements performed with the same setup give an estimate of the signal attenuation at which the noise level is eventually reached [131]. We find that starting from a typical pulse energy of ~ 1 nJ at a probe wavelength, an upconverted signal can still be detected after a damping of 8 decades, which already takes into account the fraction of the probe beam that is lost by diffuse reflection and the limited solid angle subtended by the collection optics. In addition, several other parameters can be adjusted to further enhance sensitivity, such as increasing the integration time of the detector, using a non-linear crystal with a higher upconversion efficiency and a larger angular acceptance (such as bismuth triborate) or increasing the fluence on the crystal. To this purpose, it is particularly convenient to increase the gate beam intensity as much as possible, since it does not interact with the sample and therefore does not present any alteration or damaging risk (in comparison, the crystal has a damage threshold of ~ 100 MW mm $^{-1}$).

” *Few things are harder to put up with than the annoyance of a good example.*

— M. Twain
(writer)

4.1 Validation of the technique

4.1.1 Sample preparation

In order to validate the quantitative accuracy of our newly developed setup, proper test samples of known optical properties must be prepared and measured. To allow a more straightforward comparison with the theoretical models, the samples should resemble as closely as possible the ideal infinitely extended, plane-parallel homogeneous slab. The samples that we prepared are made of a commercial UV-curing acrylate optical adhesive (Norland® 65) mixed with a dispersion of rutile nanoparticles with a diameter of 280 nm (Huntsman’s Tioxide® R-XL). The mixture of polymer and nanoparticles is rendered homogeneous through magnetic stirring (~ 10 min) and a ultrasonic bath (~ 1 h) kept at a temperature of 50°C . The liquid, opaque paste is then molded into a plane-parallel slab by letting it infiltrate inside an air gap of controlled thickness formed between two microscope glasses (Figure 4.1a). The two glasses are firmly held apart by micro-spherical spacers of calibrated size. Finally, the sample is exposed from both sides to UV light at 365 nm for a time ranging from few minutes to few hours depending on the thickness of the sample. In fact, TiO_2 nanoparticles strongly absorb in the UV range, which can significantly slow down the curing process. This preparation technique offers several advantages. Firstly, both the TiO_2 nanoparticles and the optical polymer are well-known optical materials that have been characterized thoroughly in the literature. At the probe wavelength of interest, both materials present a vanishing absorption coefficient. By controlling the density of the nanoparticles in the mixture one can therefore produce samples with controlled scattering strength. As long as the volume fraction occupied by the scatterers is kept low ($\sim 1\%$), the scattering strength can be calculated using Mie theory or numerical methods, thus providing an independent validation. Previous characterization of the same particles returned a scattering anisotropy factor of $g = 0.6$ and a scattering cross section of $\sigma_s \sim 0.1 \mu\text{m}^2$ [102, 129] for the nanoparticles, and a refractive index of 1.514 for the polymer [141], evaluated at probe wavelength.




Fig. 4.1.: (a) Top and side view of typical glass-enclosed samples. The glass slides are glued together at a controlled distance using the Norland adhesive (transparent regions). Successively, the opaque paste is sucked by capillary forces inside the central air gap and UV-cured. (b) By spin-coating in advance the glass slides, they can be later removed to get a free-standing scattering sample. (c) Cross-section of the free-standing slab sample observed at the scanning electron microscope. The spatial distribution of scattering particles shows small-scale clusters homogeneously dispersed across the sample thickness. As confirmed by our experimental characterization and by the literature on TiO_2 nanoparticles [142], light propagates perceiving an effective, homogeneous scatterer density even in the presence of small local clusters.

From an experimental point of view, a particularly challenging issue that was addressed during this work is that of boundary conditions. The samples prepared with a previously used procedure, in fact, cannot be removed from the enclosing glass slides. In principle, this configuration is compatible with a diffusion-based model with matched boundary conditions (the polymer has roughly the same refractive index of the glass slides), but in practice the finite thickness of the enclosing slides makes data evaluation more complex. In fact, the sample is effectively a 3-layer slab embedded in air, which causes a large fraction of light being reflected back into the sample from air-glass boundaries. Using 1 mm-thick slides, the reflection-free time window is defined by the ballistic time that diffused light takes to be reflected back into the sample $2 \times 1 \text{ mm} / v_{\text{glass}} \sim 10 \text{ ps}$, meaning that all spatial and temporal information collected after this delay will be deformed by the presence of this highly convoluted extra contribution. While a 3-layer configuration could be easily accounted for in MC simulation, data evaluation with the diffusive approximation on a time scale longer than few ps becomes much less straightforward.

Few ways of reducing this problem have been attempted in the past, such as gluing thicker glass layers of even absorbing filters to suppress internal reflections [129]. Eventually, a new sample fabrication technique has been developed, allowing to produce free-standing scattering samples with well-defined flat interfaces. By using a spin coater, a controlled sub-μm layer of polyvinyl alcohol (PVA) is deposited on both glass slides before they are glued together and infiltrated with the opaque paste. Polyvinyl alcohol is a water-soluble polymer, and can be dissolved after curing the sample to release it from the glass slides (Figure 4.1b). Figure 4.1c shows a cross-section of the turbid region of the sample. As can be qualitatively seen, the TiO_2 nanoparticles are dispersed in small clusters homogeneously distributed across the thickness of the sample. It has been

experimentally demonstrated that this slight level of inter-particle clustering does not differ substantially from a perfectly homogeneously dispersed mixture [142]. Our optical characterization, as described in the following, confirms that light perceives an effective, homogeneous scatterer density even in the presence of small clusters.

4.1.2 Data evaluation models

The second critical element in our validation procedure is represented by the correct implementation of the fitting models. In our investigation, the experimental data is compared both with the diffusion approximation and with the Monte Carlo model of light transport. In the latter case, due to the flexibility of the simulation software, an accurate configuration can be implemented, including the actual spatial and temporal distribution of the probe beam. As regards the diffusive approximation, care has to be taken when adapting the actual experimental configuration to the simplifications of the model. Few points in particular are worth discussing. As we discussed in Section 2.2.3, the relations that we obtained for the transmittance from a slab refer to an isotropic point source situated at a depth z_{src} rather than to a pencil beam emitting a δ -like pulse at $t = 0$ s. In order to use a relation for the total transmittance such as equation (2.113), the experimental time scale of the measurement must be translated integrally to match the origin of the time axis. Determining the correct time origin is a critical task which can affect dramatically the accuracy of an optical characterization technique. If the time origin is set too early or too late, unphysical results might arise (such as pre-ballistic transmission), and the characterization of the scattering properties would be heavily affected because even a slightly anticipated or delayed rising time in the transmittance curves can be associated to higher or lower scattering probabilities.

The steps involved in determining the correct origin of the time axis are depicted in Figure 4.2 for a typical sample. The starting reference is of course provided by a cross-correlation measurement taken without any intervening sample, which will reach its maximum at a given position of the translation stage (Figure 4.2a). This delay does not represent yet the origin of the time axis because, even in the cross-correlation case, the probe pulse arrives at the crystal after passing through an effective ‘sample’ made of air with the same thickness L as the actual sample. By knowing in advance the thickness of the sample, it is therefore possible to anticipate the cross-correlation signal (or, equivalently, delay the transmittance) by an offset of L/c to account for it (Figure 4.2c). To our knowledge, this is the only correction, if any, that is usually applied in the literature in order to set the time axis. However, a finer adjustment is still required to match the assumption of the DA model, relating to the fact that the pulse should be emitted from *within* the sample, not from its outer surface (Figure 4.2d). As a matter of fact, the DA requires that the isotropic source is placed at a depth $z_{\text{src}} = l'_s$. Clearly, this extra step compromises any optical characterization procedure in the first place: strictly speaking, it is not even possible to set the time axis of the measurement without knowing in advance the scattering mean free path. To circumvent this issue, in our analysis we developed a new iterative fitting technique for the retrieval of the optical properties. Starting from a reasonable guess for l'_s , the time axis is iteratively shifted




Fig. 4.2.: Procedure followed to match the origin of the absolute time scale to that of the isotropic point source of the DA model. (a) Raw output for the total transmittance and cross-correlation counts of a typical sample as a function of the translation stage position. (b) Conversion to relative delay using the relation $\Delta t = 2\Delta z/c$. (c) Correction of the relative delay by the time taken by the probe pulse to cross the empty region corresponding to the sample thickness. (d) Correction of the relative delay by the depth at which the DA assumes to have an isotropic point source. (e) Final time shift of both measurements to set time origin at the obtained delay for the cross-correlation.

by l'_s/c until the result of this self-consistent fit converge to the previous guess within a fixed tolerance. By doing so, variations of a few percent are to be found with respect to the previous procedure even for samples with an optical thickness $\gtrsim 20$, well in the diffusive regime. Still, any procedure devised to evaluate the total transmittance curve necessarily requires the knowledge of the effective refractive index, refractive index contrast and exact thickness of the slab. These parameters are often not known to a high precision, or maybe not even well defined when the sample presents a more complex heterogeneous structure.

In the case of our reference samples, determining the thickness and the refractive index can be conveniently performed on an all-optical basis. By detecting both the main cross-correlation signal and its first replica corresponding to an additional round-trip inside the slab, both the refractive index of the polymer and the thickness of the slab can be retrieved. The former can be estimated by comparing the ratio of the integral of the two peaks, which are proportional to $(1-r)^2$ and $(1-r)^2 r^2$, with $r = \left(\frac{1-n}{1+n}\right)^2$ assuming normal incidence, while the latter is retrieved directly by the raw distance of the two peaks as $L = \Delta z/n$. The obtained values are in excellent agreement with the specs of the polymer and the scanning electron microscope (SEM) images of the sample (Figure 4.1c).

The second main data evaluation model that we will use in this chapter is represented by Monte Carlo simulations, which we perform using a new C++ software library called MCPLUSPLUS [143] that was developed from scratch during this thesis work in close collaboration with Dr. G. Mazzamuto (see also Chapter 5). Similarly to other available software packages [53, 144], MCPLUSPLUS allows to model a sample as a stack of layers of different properties. Additionally, few others unique features are available which are dedicated to an accurate modeling of the light source. Unlike other available libraries,




Fig. 4.3.: Example of Gaussian ray bundle source with a beam waist of $8\text{ }\mu\text{m}$, focused in $(0, 0, 0)$. At each cross-section along z the intensity is gaussian.

our software offers several angular, temporal and spatial distributions to choose from as regards the source term. In particular, a stochastic $\text{sech}^2(t)$ pulse shape model is available exploiting the analytical inverse of its cumulative distribution [145], which allows to model with extreme accuracy the temporal evolution of the probe pulse and the instrument response. Analogously, among several simpler models, MCPLUSPLUS offers a simple and extremely efficient algorithm to model accurately a gaussian beam, which has been developed following reference [146]. Figure 4.3 shows an example of the trajectories that form a gaussian ray bundle focused on the entrance surface of the sample. The model allows to select the waist of the beam at the focusing lens, the waist in the focus, and the working distance of the lens itself. The resulting beam satisfies the correct gaussian intensity profile at each cross-cut along z . As a further validation of the validity of this approach, propagating the beam through several non-scattering layers of different refractive indexes, results in the correct refocusing distance and waist by simply applying Snell's law or ray transfer matrices to each of the trajectories. Therefore, this source model can also be used for the case of samples enclosed in glass slides or containers, assuming that, as it is commonly the case, the probe beam is focused at the interface of the sample rather than of the container.

If not stated otherwise, all simulations performed to describe our scattering samples have been performed assuming a gaussian-ray-bundle source term with a spatial full-width-half-maximum of $10\text{ }\mu\text{m}$ and a temporal width of $\Delta\tau_p = 170\text{ fs}$. For the polymer samples, a scattering anisotropy of $g = 0.6$ has been assumed, as calculated for a 280 nm -sized TiO_2 nanoparticle using Mie theory [26].

4.2 Unveiling data evaluation artifacts

4.2.1 Limited thickness vs absorption

The first set of measurements that we present refers to the simplest sample possible, i.e., a plane parallel homogeneous scattering slab as depicted in Figures 4.1b and 4.1c. The

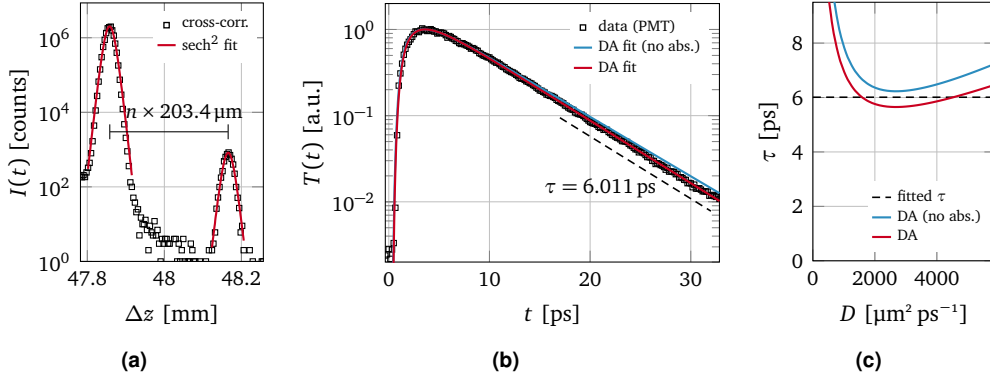


Fig. 4.4.: Set of measurements on the reference sample. (a) Optical measurement of the thickness and refractive index of the polymer slab. The delay between the main peak and its internally reflected replica gives directly nL , while the ratio between the integral of the two pulses depends only on n . We obtain a thickness of the slab of $203.4 \mu\text{m}$, in perfect agreement with the SEM image of Figure 4.1c. (b) Experimental transmittance data (black squares) acquired with the PMT compared to an exponential fit of the asymptotic decay time (dashed line). Two different DA fits are shown, forcing null absorption and leaving it as a free parameter. As shown in panel (c), the obtained decay time is not compatible with a non-absorbing sample according to the DA.

initial characterization consist of measuring the total transmittance using the PMT as an integrated photon counter. A fit of the experimental data point is performed using equation (2.113) to retrieve the scattering properties of the slab. In the fitting routine, relative rather than absolute residuals are minimized, in order to weigh properly both the early part of the curve and its asymptotic decay. As discussed previously, the experimental data can be normalized since all the information can be retrieved from the shape of the curve rather than from absolute fluence levels. In principle, we expect that the sample is purely scattering, with no absorption. Yet, a DA fit with $\mu_a = 0 \text{ mm}^{-1}$ gives actually a poor agreement with the data, as we are able to appreciate thanks to the high dynamic range of the detector (Figure 4.4b). It is worth noting that this failure cannot be attributed to the selection of a non-optimal fitting range (which can potentially affect the results of the fit to a substantial extent [147]). In fact, as Figure 4.4c clearly demonstrates, the experimental asymptotic decay time is altogether incompatible with the DA, unless some absorption is included. Conversely, a perfect agreement between the model and the data is obtained with a two-parameter fit, returning $l'_s = (24.1 \pm 0.6) \mu\text{m}$ and $l_a = (12.0 \pm 1.1) \text{mm}$, corresponding to an albedo of $\varpi_0 = 0.998$.

A few comments on the reliability of the obtained results are in order. The value of the absorption coefficient returned by the fitting routine is indeed very small, as underlined by the high albedo. The predicted optical thickness of the slab would be of ~ 8.4 , with 8 being usually considered the rule-of-thumb lower limit above which the DA can be regarded as an acceptable approximation [72]. Yet, the total integrated absorbed energy associated to the slab would be equal to $A = 1 - T - R \simeq 8\%$, which is definitely higher than expected for this sample. We should note that the apparent inconsistency of the results arises from the fact we are testing a reference sample of




Fig. 4.5.: Set of time-resolved frames acquired at different delays with our ultrafast imaging setup. Each frame is averaged over different disorder realizations (different regions of the sample) and is displayed normalized to its own maximum intensity.

known properties, while it would be unreasonable to question a similar measurement performed on a truly unknown medium, given the exceptional agreement of the fit, the high albedo and acceptable optical thickness.

The inconsistency of the retrieved optical parameters can be completely resolved by means of a full spatio-temporal investigation and non-approximated modeling. In fact, the analysis of the time evolution of the mean square width (MSW) of the transmitted profile is particularly interesting in that it is inherently free from this absorption-to-scattering crosstalk effect [110, 148, 149]. The MSW, defined as the variance of the spatial profile (2.103), can be directly determined if the full spatio-temporal evolution of the $T(\rho, t)$ (or $R(\rho, t)$) profile is known. Due to the inherently normalized definition of the MSW, any amplitude factor (such as absorption) applied to the profile will cancel out exactly at any time, leaving its shape unmodified (see Figure 2.7). We therefore further investigated the same sample recording a set of transmittance profiles at different delays with the setup in the imaging configuration (Figure 4.5). As can be seen, the experimental $T(\rho, t)$ profiles broaden in time as predicted by the diffusive approximation. The time-dependent values of the variance can be extracted assuming a diffusive profile shape (gaussian), or by directly calculating it through its general definition (2.103) which does not rely on any profile-shape hypothesis. The linear growth of the MSW is plotted in Figure 4.6a, and can easily be interpreted within the diffusive approximation using its characteristic prediction $w^2(t) = 4Dt$ (2.104).

The slope retrieved experimentally corresponds to a $D_{\text{exp}} = (1746 \pm 21) \mu\text{m}^2 \text{ps}^{-1} \rightarrow l'_s = (26.6 \pm 0.3) \mu\text{m}$, which is appreciably larger than the value retrieved from the time-resolved curve (Figure 4.4b). In order to resolve this discrepancy, we resort to a full MC modeling of the problem. By performing a MC fit of both the $w^2(t)$ and $T(t)$ curves, we eventually succeeded in simultaneously reproducing both experimental data sets using a single transport parameter l'_s (Figure 4.6). The numerical inversion procedure




Fig. 4.6.: (a) Linear growth of the experimental mean square width of the profiles (circles), compared with the result of a MC simulation with $l'_s = 25.5 \mu\text{m}$ and $\mu_a = 0 \text{ mm}^{-1}$ (solid line). A linear fit of the experimental data (not shown) returns a value of $l'_s = 26.6 \mu\text{m}$ according to equation (2.104). (b) Experimental transmittance data (black squares) acquired with the PMT compared with the result of the same non-absorbing MC simulation. The integrated intensity of the UFI frames is plotted as empty circles, showing the excellent linearity of the CCD detector as compared to the PMT.

returns a value of $l'_s = (25.5 \pm 0.5) \mu\text{m}$, which is perfectly compatible with both curves without the need to add any absorption contribution.

It should be noted that, as compared to MC simulations, the two independent DA-based estimates of the scattering properties that we obtained return both a wrong value, but by a different amount and for very different reasons. As regards the transmittance decay, the result is particularly far from the actual value both because of the intermediate level of turbidity and due to the cross-talk effect with absorption, which overcompensates the failure of the DA to some (uncontrolled) extent. On the contrary, the reduced scattering mean free path obtained from the MSW slope results in just a 4% overestimation of the actual best MC estimate. A comparable level of accuracy is acceptable for a wide array of applications, where ultrafast imaging (UFI) techniques could be exploited to extend the applicability of the DA in this intermediate thickness range. Moreover, as widely suggested in the literature [148, 149], estimating transport properties from the mean square width rather than from time-resolved data is much more accurate and straightforward because in the former case, *a priori* knowledge of the absorption coefficient, refractive index contrast and sample thickness is not required. Likely, the DA holds more robustly against a high refractive index contrast and a low optical thickness for in-plane transport, because effects coming from thickness truncation are less relevant for light propagating along the slab rather than perpendicularly to it. This reflects in the fact that the simple $4Dt$ dependence does not even involve any parameter other than the diffusion constant, including the geometry of the sample or its extrapolated boundary conditions. Most importantly, the simple determination of D enabled by the measurement of $w^2(t)$ is also insensitive to the exact determination of the time axis origin and the exact time-window considered for the fitting, both of which represent long-standing issues in the evaluation of time-resolved data [147, 150], as discussed previously.




Fig. 4.7.: (a) DA-based analysis of the total time-resolved transmittance measured on the second sample. An excellent agreement is obtained only with a two-parameter fit including absorption. (b) Linear growth of the MSW as retrieved from the frames of Figure 4.8. A linear fit is shown as a guide for the eye.

In addition, for the same reason that the MSW is independent of the absorption, it is also effectively independent of the integration time, as well as of any laser power fluctuations, which makes the technique inherently more robust to drifts and variations in the experimental conditions. Moreover, as shown in Figure 4.6b, our UFI technique is capable of replacing completely integrated measurements with comparable sensitivity. In this respect, a single measurement session is capable of providing two almost orthogonal data sets: indeed, the former does not depend on the integrated intensity of the individual frames, while the latter is primarily linked to the transport along the z -axis. The independence of $w^2(t)$ from the integration time represents a strong advantage from an experimental point of view, allowing to optimize the exposure time on a per-frame basis. As regards the frames of Figure 4.5, each frame has been integrated for a time of the order of a minute, which was kept constant in this circumstance to allow a direct comparison also of the integrated intensity so to verify the linearity of the CCD. Nonetheless, the fluence transmitted by the sample at this optical thickness was entirely sufficient to cover a remarkable time span of roughly 5τ . Our perfectly matching Monte Carlo validation, which is conducted independently on two experimental curves that depend differently on different sets of parameters, is a clear indication of the quantitative accuracy and reliability of our setup over many orders of magnitude in both the temporal and spatial domains.

4.2.2 Layered heterogeneities vs anisotropic transport

A second relevant sample that we analyzed with our combined UFI and total transmittance description is a new polymer slab with similar thickness ($190\text{ }\mu\text{m}$) and doubled scatterer density. As opposed to the previous sample, this slab falls well within the diffusive regime in terms of its final turbidity. Also in this case, a DA fit of the time-resolved curve yielded a perfect agreement (Figure 4.7) only when accounting for the presence of a small but unexpected amount of absorption. The values returned by the fitting routine




Fig. 4.8.: Set of time-resolved frames acquired for the second diffusive sample. Each frame is averaged over different disorder realizations (different regions of the sample) and is displayed normalized to its own maximum intensity.




Fig. 4.9.: Comparison of experimental data with two MC simulations matching, respectively, the experimental asymptotic decay time (dash-dotted line) and the MSW growth rate (dashed line). The two simulations return $l'_s = 14.2 \mu\text{m}$ and $l'_s = 15.7 \mu\text{m}$. (a) Both simulations reproduce poorly the experimental time-resolved data, as can be better appreciated in the linear-scale inset. (b) The high in-plane diffusion rate can be matched using a long mean free path close to that found with the simple DA prediction.

are $l'_s = (12.0 \pm 0.4) \mu\text{m}$ and $l_a = (13.2 \pm 1.2) \text{mm}$, which together determine a total absorption of $A \sim 7\%$ despite the extremely high albedo.

From the analysis of the UFI frames (Figure 4.8), a perfectly linear MSW growth is obtained, as shown in Figure 4.7b. Nonetheless, the retrieved slope of $D_{\text{exp}} = (1050 \pm 13) \mu\text{m}^2 \text{ps}^{-1}$ is much steeper than expected from the time-resolved data, yielding a l'_s of $(16.0 \pm 0.2) \mu\text{m}$ as if the diffusion process was enhanced along the in-plane directions.

In sharp contrast with the previous experiment, in this case a non-absorbing Monte Carlo fit was unable to reproduce the time-resolved data, even when the observed experimental decay time or the MSW slope were perfectly matched (Figure 4.9).

Driven by this discrepancy and by the characterization capabilities enabled by the UFI setup, we inspected the cross-section of this sample under optical and electronic microscopy. This revealed a heavily layered modulation of the scatterer density compared




Fig. 4.10.: (a) Montage of optical and SEM images of the second, heterogeneously disordered sample. The central inset shows a transverse section, which exhibits a layered symmetric structure that we modeled with a high-density central layer ($L_1 = 26 \mu\text{m}$), a low-density interstitial layer ($L_2 = 60 \mu\text{m}$) and an intermediate-density outer layer ($L_3 = 22 \mu\text{m}$). The heterogeneity represented by the core layer is also appreciable at the optical microscope (left inset). The right insets highlight the density difference between regions 1 and 2. (b) Lateral section of third investigated sample, a specimen of dried integument of a grapefruit segment.

to the more homogeneously dispersed one of our first sample (compare Figures 4.10a and 4.1c).

Taking advantage of the versatility of MC simulations, we modeled our sample after the SEM images as being composed of five layers of three different thicknesses and densities, arranged symmetrically with respect to the central layer. A new MC brute-force fit assuming a fixed geometry and only $l'_{s,1}$, $l'_{s,2}$ and $l'_{s,3}$ as free parameters, was eventually able to perfectly reproduce both the time-resolved and MSW curves simultaneously as shown in Figure 4.11. The absorption coefficients were all set to $\mu_a = 0 \text{ mm}^{-1}$. We obtained mean free paths of $l'_{s,1} = (3.5 \pm 0.5) \mu\text{m}$, $l'_{s,2} = (21.5 \pm 0.5) \mu\text{m}$ and $l'_{s,3} = (11.0 \pm 0.5) \mu\text{m}$, respectively, for the high, low and intermediate-density layers. In this case, the combined MSW/time-resolved decomposition clearly allowed us to reveal an unexpected degree of complexity, signaling the inconsistency of the homogeneity assumption. Together, the two data sets cast a rigid constraint on the inverse problem solution, involving both longitudinal and transverse propagation dynamics. The analysis that we performed assumed a known structure, but there are indeed many practical circumstances where the geometry of the problem is known, and the investigation is focused in retrieving the optical properties of each layer. Moreover, layered structures can often be associated with anisotropic transport, and vice versa. In this sample, we could study the effect of a layered structure separately from that of anisotropic transport, since a homogeneous pattern of unconnected point scatterers is inherently isotropic.

Interestingly, resolving this composite structure showed that a layered heterogeneity can, in principle, mimic the effects of anisotropic transport which, to our knowledge, is an effect still unaddressed in the literature. This is especially relevant given the pervasiveness of layered media (for example in coatings industry, atmospheric physics and biological tissues), which often exhibit counter-intuitive features that standard




Fig. 4.11.: Comparison of experimental data with a 5-layer MC simulation modeled after Figure 4.10a. The fit returns the scattering properties associated to the 3 different densities of scatterers, in good accordance with a qualitative inspection of the cross-section of the sample. Panel (a) shows the superposition of integrated intensities obtained with the PMT and CCD camera.

time-resolved techniques are still unable to explain, as in the case of light transport across the human forehead [151].

4.3 A novel transport regime in ultra-thin samples

In the previous section, we have explored the advantages offered by a decomposition of the transmittance data into a total time-resolved curve and a MSW growth. Nonetheless, it is clear that the raw, non-integrated frames provide an irreducible set of information. A dramatic illustration of this point is obtained by studying a thin biological sample, such as a small strip cut from the dried skin of a slice of grapefruit. As the SEM image reveals (Figure 4.10b), its structure consists of a conglomerate of small flakes forming a corrugated slab $\sim 85 \mu\text{m}$ -thick. The dried sample appears to be a brittle, almost transparent membrane, as shown in Figure 4.12a.

In this range of optical (and absolute) thicknesses, standard experimental techniques fall short because of the extreme time scales involved, causing common diffusive modeling to fail drastically. Furthermore, most of the signal will be ballistically transmitted, carrying almost no information regarding the optical properties of the sample. As Figure 4.12b shows, the fraction of light that is scattered inside the grapefruit membrane is less than 1/50th of the original intensity. Nonetheless, our ultrafast imaging technique allows us to selectively address the light that was retained for a longer time inside the sample whereas at the same time spatially inspecting it as it propagates along the main (transverse) slab extension. Figure 4.13 shows a collection of frames acquired over a time-window of ~ 5 ps after pulse injection, corresponding to a total path length greater than 13 times the sample thickness. We see that the light spreads through the membrane with a well-defined wavefront traveling inside its main plane, resulting in a dramatic departure from any prediction compatible with the diffusive framework. However, standard time-resolved or steady-state investigations would still just measure a

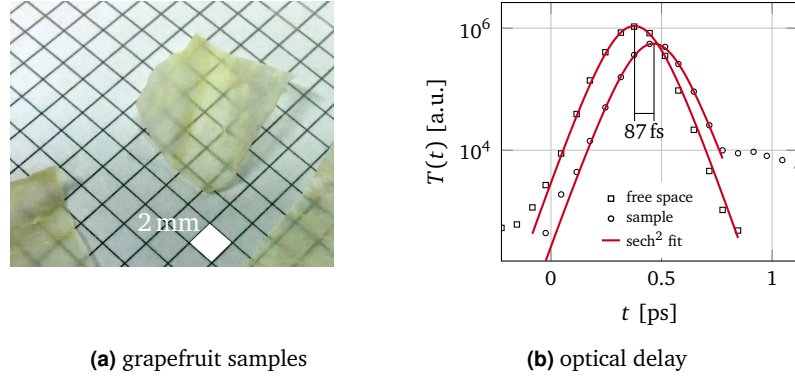


Fig. 4.12.: (a) Picture of the grapefruit samples. The membranes have an almost transparent appearance. (b) Optical delay introduced by the presence of the grapefruit membrane on the probe beam path. By knowing its thickness from previous SEM images (Figure 4.10b), the effective refractive index is found as $n = 1 + c\Delta t/L \sim 1.31$.




Fig. 4.13.: Set of time-resolved frames acquired for the weakly scattering biological sample. Each frame is averaged over different disorder realizations (different regions of the sample) and is displayed normalized to its own maximum intensity. The propagating light pattern differs dramatically from that of more turbid media, exhibiting minimum intensity along the injection axis.

single decay time and a bell-shaped profile, respectively, which could deceptively support an inappropriate interpretation in terms of the DA. Additionally, the investigated sample exhibits various signatures of anisotropic light transport, with the luminous wavefront propagating differently along the x and y direction.

As we will show, the UFI experiment provides a deep insight into the light transport properties of complex systems, which requires the development of a novel analysis methodology. As a proof-of-concept, we demonstrate a procedure to assess in-plane transport properties in this extremely optically thin regime, an experimental task for which there are basically no available characterization techniques up to date. For the sake of simplicity, we will illustrate our investigation on a cross-cut of the profile assuming isotropic transport. The general treatment is analogous and requires the use of an anisotropic transport model (not to be confused with anisotropic scattering), such as anisotropic MC simulations where $D \rightarrow \mathbf{D}$ is a tensor with different components along x , y and z [152]. Figure 4.14a shows a comparison between experimental cross-cuts at different times and the corresponding time evolution of an isotropic Monte Carlo




Fig. 4.14.: (a) Horizontal cross-sections of the wavefront and a MC fit with g and l_s as free parameters. The fit returns values of $g = 0.7$ and $l_s = 150 \mu\text{m}$. (b) Simulated radial intensity profiles in this extremely optically thin regime exhibit a significant breakdown of the similarity relation.

simulation that succeeds in reproducing the wavefront speed, intensity decay and overall shape at all given times. In order to achieve this collective agreement at each different time, the anisotropic scattering coefficient g needs to be considered as an additional parameter for the MC fitting procedure. Indeed, as shown by Figure 4.14b, the measured intensity patterns exhibit a significant breakdown of the typical DA degeneracy condition expressed by the similarity relation (2.38), with g playing a substantial role in determining the overall shape and time evolution of the traveling wavefront, even at fixed l'_s . In particular, while traveling outwards, the instantaneous position and peak intensity of the density wave vary appreciably with different combinations of g and l_s , therefore allowing both to be retrieved within an error of a few percent points. The set of simulated curves shown in Figure 4.14a is obtained using $g = 0.7$ and $l_s = 150 \mu\text{m}$, corresponding to an in-plane $l'_s = 500 \mu\text{m}$. Qualitatively accurate figures for this particular specimen would require a numerical analysis involving fully anisotropic simulations, yet the aim of this experiment was to demonstrate the existence and utility of this peculiar in-plane transport regime and its onset in optically thin media.

This experiment showcases the characterization potential enabled by a rich time- and space-resolved output. Our estimation of g results from a collective fitting involving multiple spatio-temporal profiles, standing in contrast with more common and less robust characterization techniques involving a single scalar measurement, such as the attenuation coefficient of a collimated beam [92, 153]. Accurate determination of single scattering anisotropy is crucial to directly probe the shape, orientation and optical properties of single microscopic scatterers and might therefore be of interest to a broad range of fields, from realistic computer rendering of participating media [16, 154] to photo-therapeutic diagnostics of biological tissues [155, 156]. Moreover, compared to other available techniques, the determination of g and l_s is achieved in a single measurement session, which also simplifies the retrieval procedure. The peculiar wave-shaped intensity pattern that enables such a rich characterization of the scattering properties is of course not unique to this particular sample, and we observed it in a




Fig. 4.15.: Examples of other specimens exhibiting a similar in-plane transport pattern as shown by the grapefruit sample. (a)-(b) Dried onion skin layer with a thickness a few μm . The elongated shape and alignment of the plant cells induce an anisotropic propagation of the luminous pattern (frames are recorded with a $\Delta t = 0.67$ ps). (c)-(d) Wing of a *Pieris rapae* butterfly ($\Delta t = 0.33$ ps).

variety of different biological samples of both plant and animal origin, usually differing only by the actual contour shape of the propagating wavefront. Figure 4.15 shows qualitatively two such examples obtained for a dried onion skin specimen and the wing of a *Pieris rapae* butterfly. This supports the general validity and usefulness of our proof-of-concept evaluation technique in studying extremely complex media, included biological tissues.

4.3.1 Outlook and open questions in the physics of light transport


To summarize, the optically-gated imaging technique that we developed offers several advantages, such as a high qualitative fidelity over a large field of view with respect to Fourier-plane configurations, and sub-ps time resolution compared to other electronically-gated imaging solutions. The applicability of our technique is very general and non-invasive, usable at different wavelengths, with a broad field of view and high temporal resolution. We therefore envision that it could be applied to a wide range of photonic applications for both the sub-ps physics it allows to investigate and its convenient wide-field acquisition, which does not require scanning over the region of interest. In particular, gaining direct access to such time scales not only enables the study of optical properties in thin or inherently minute specimens (e.g., the ocular fundus, vascular walls, skin dermis or dental enamel), but is fundamental to investigate light transport at the mesoscopic level.

In this respect, the recent emergence of experimental alternatives to harness the full potential of a complete spatio-temporal characterization of light transport [42, 110, 157], holds promise for a new generation of experiments capable of tackling the many open questions that are object of debate in the literature of light transport.


As we demonstrated in the previous subsections, gaining access to transverse and axial transport combined allows to evaluate correctly the applicability of the common diffusion framework to avoid the subtle pitfalls and artifacts that can arise under certain circumstances. In particular, the direct measurement of the transient MSW expansion enables a robust and simple interpretation. In sharp contrast to other common observables, such as the total transmittance and its decay time, the MSW is independent of absorption and largely unaffected by refractive index contrasts and sample thickness, therefore overcoming long-standing problems posed by their precise assessment. One prominent case where this feature could be decisive is that of the experimentally observed decrease in the diffusive coefficient of a turbid slab with decreasing thickness [87], which was indirectly determined through the transmittance decay time. A spatio-temporal measurement in terms of the mean square width evolution would directly probe the diffusion coefficient irrespective of the exact boundary conditions, the incorrect estimation of which has been cited as a possible cause for this apparently anomalous behavior [71, 158].

Analogously, the study of structurally anisotropic media would be largely facilitated by the three-dimensional characterization capabilities of a spatio-temporal investigation. This class of samples is extremely important in many industrial and biomedical applications, yet their theoretical modeling in terms of light transport is still object of debate in the literature [152, 159]. Further interest arises from the fact that anisotropic transport seems to represent a key factor in enhancing the scattering strength of a material beyond expectations [131, 160]. As a matter of fact, anisotropic transport modeling is only one examples of numerical effects still lacking a thorough experimental validation. Other interesting topics that might be investigated with novel techniques include the polarization dependence predicted for the diffusion coefficient [45], as well as the occurrence of a non-monotonic MSW evolution in samples exhibiting Anderson localization [149] or the lateral expansion of Lévy-type transport, where the experimental limit posed by the finite thickness of the sample might be circumvented [29]. Similarly, an UFI measurement of intensity profiles reflected by macro-porous bulk materials could access light spreading within the first few hundred fs, from which the time-dependent diffusion constant could be determined [161], leading to information concerning the structure factor of the material hardly obtainable at long times.


The wide-field parallel acquisition demonstrated by the setup might, however, find applications in many different fields other than the study of disordered media and their transport regimes. One prominent example is that of the characterization of waveguide devices, which are typically investigated using reflectometry methods, near-field scanning optical microscopy or atomic-force microscopy. More recently, an ultrafast time-resolved technique has also been proposed, allowing to address individual optical components *in situ*, which is desirable with respect to other methods requiring additional fabrication steps or test structures [162]. Still, to date, the study of single integrated optic elements is performed by scanning point-by-point over the region of interest of the sample, which is extremely time consuming. Conversely, taking advantage of our UFI technique, it is possible to inspect and measure directly the light that is coupled out of the device (e.g., from a defect) over a wide field-of-view. Figure 4.16 shows an illustrative measurement



(a) polymer waveguide coupled to a ring resonator (SEM picture)



(b)



(c)

Fig. 4.16.: Application of the UFI technique to the investigation of polymer photonic circuits. (a) The system is composed by a waveguide vertically coupled to a ring resonator. Input and output gratings are optimized for perpendicular injection and extraction of light. (b) UFI signal received from the output grating. A direct-illumination image of the sample is superimposed as a reference. (c) Time traces measured by spatially selecting the input and output gratings, respectively. A long series of time-resolved replicas are emitted from the output grating, revealing the complex interplay between the waveguide and the ring resonator.

performed on a simple optical circuit made of a polymer waveguide and a polymer ring resonator vertically coupled together on a glass substrate. The samples have been fabricated using direct laser writing by S. Nocentini.

Given the typical applications of optical circuits for telecommunications, the system is optimized for a wavelength of 1550 nm. In the measurement (performed by D. Nuzhdin) we have therefore exchanged the probe and gate beam, slightly changing the wavelength of the latter to match the nominal working wavelength of the device. Figure 4.16b shows a typical UFI frame acquired at the delay of maximum output intensity (direct-illumination image of the sample is superimposed as a reference to recognize the output grating). However, a much richer dynamics develops at later time, signaling the presence of the ring resonator and its interplay with the waveguide. Measuring the signal recollected from the gratings, or even the light that is leaked by other optical elements or unintentional defects, enables to assess the correct functioning of the device and its specifications from the global to the local scale. Further study on this topic is needed to better understand the full potential of the technique and the role of its narrow optical acceptance with respect to this kind of applications.

Space-time characterization of the ballistic-to-diffusive transition: application to the inverse problem

” *An approximate answer to the right problem is worth a good deal more than an exact answer to an approximate problem.*

— J. Tukey
(mathematician)

5.1 Inverting light transport in a scattering slab

As we discussed in Chapter 2, the forward problem of light propagation in turbid media is highly nonlinear. This means that the inverse problem cannot be solved by simple inversion without introducing approximations. The most widely adopted approach is therefore iterative. With reference to the general description introduced in Chapter 1, the inverse model consists in retrieving the set of optical properties $\mathbf{p}(\mathbf{r})$ given that we have measured the reemitted light $X_{\text{out}}(\mathbf{r}, t, \mathbf{s})$, or some portion of it. In a typical iterative approach, we start by guessing initial values of the optical properties $\mathbf{p}(\mathbf{r})$ and feeding them into the forward model of choice. The computed values of X_{out} are then compared with those measured, and using a suitable minimization algorithm the values of $\mathbf{p}(\mathbf{r})$ are updated. The process is repeated until the computed and measured values coincide within the target accuracy.

In this Chapter we present an extensive characterization of the delicate transition between the ballistic and the diffusive transport regime. Taking advantage of the obtained results, we will discuss some important aspects of the inverse problem, and describe a new non-iterative method to retrieve the optical properties of an unknown medium based on robust, spatio-temporal descriptors. The first important classification for inverse problems concerns the number of unknowns, which is often given by the number of distinct spatial regions multiplied by the number of optical properties used in the model. An important case at the core of many applications is that of a homogeneous and isotropic medium, where the number of unknowns $\mathbf{p}(\mathbf{r})$ is reduced to the number of optical parameters. Secondly, an appropriate forward model and type of measurements should be selected. Ideally, the number of observables should be equal to the number of unknown parameters, and they should all be orthogonal in the measurement space.

The iterative nature of the common data evaluation process poses a set of strict constraints on the forward model, whose evaluation time must be short in order to reach convergence in a reasonable time. While the analytical form of the diffusion model is attractive and straightforward to evaluate, its underlying assumptions break down in many relevant experimental circumstances and should be therefore avoided whenever possible. More recently, analytical solutions to the radiative transfer equation are also being developed which could streamline the inversion procedure, but are still limited to the steady-state domain [50], or to infinite and semi-infinite media [51, 163]. Moreover, numerical stability and performance issues still exist when manipulating these solutions, hindering their application to the slab geometry [164]. On the other hand, the gold standard of Monte Carlo (MC) simulations represents an ideal forward model in terms of accuracy, yet to a higher computational cost. Nonetheless, the increasing availability of massively parallel computation capabilities is gradually fostering Monte Carlo simulations as a viable method to solve the inverse problem. To reduce the computational burden of iterative procedures, MC fitting routines to date have exploited rescaling properties of the radiative transfer equation to adapt a limited set of pre-simulated Monte Carlo data to experimental measurements [91, 147, 165–172]. In order to limit the occurrence of ‘scaling artifacts’, rescaling must be typically performed on a single-‘photon’ basis, thus requiring to store each exit time and position separately [170, 171]. Bin-positioning strategies are also known to represent a possible source of artifacts, requiring complex correction strategies to be deployed [172]. Finally, while for the semi-infinite geometry a single dimensionless Monte Carlo simulation can be rescaled both in terms of the absorption and scattering mean free paths using equations (2.52) and (2.53), the computational cost increases in the case of finite thickness geometries, where different scattering mean free paths values must be simulated separately. This is why only few examples of MC-fitting routines can be found in the literature dealing with this configuration [167].

Completely different strategies exist, though, that do not involve iterative fitting in the first place. These include for example least-squares support vector machines, neural networks, or simpler non-iterative fitting [173–177]. In addition to these, a large and extremely general group of methods is represented by lookup table (LUT) approaches, where the inverse problem for a given configuration is exhaustively solved in advance by brute-force. In the literature, several lookup-table approaches have been proposed, based on both experimental [178] and synthetic data [93, 104, 168, 179–182]. In the latter case, the LUT routine is comprised of a database of pre-simulated realizations with different optical parameters, that is then queried directly to solve the inverse problem. In sharp contrast to fitting procedures, in this case the computational effort represents a one-time cost concentrated in the compilation of the lookup table, which can be accessed instantaneously later on.

As opposed to a standard fit, lookup-table routines rely on single scalar parameters directly linked to transport properties. Typical examples are represented by the total amount of transmitted/ballistic/reflected light from a slab. This triplet of observables, often referred to as T_{tot} , T_{coll} and R_{tot} , has been extensively exploited to retrieve optical parameters through MC-LUT routines [93, 104, 179–183], despite the fact that such absolute intensity measurements are extremely challenging [93, 104] and prone to

unpredictable systematic errors [181]. Moreover, while the natural propensity of MC-based data evaluation techniques is towards the study of optically thin media, integrated transmittance/reflectance quantities do actually lose their usefulness as the sample thickness decreases [91], since the acquired signal will be eventually dominated by light that has been either specularly reflected or ballistically transmitted through the sample, thus carrying almost no information about its properties. Even more fundamentally, one of the main assumptions behind the use of integrating spheres is that of a lambertian diffuse profile [179], which is clearly not holding for thinner systems.

In the following sections, we aim at improving the common LUT approach to the inverse problem taking advantage of the robust observables made available through spatio-temporal experimental techniques such as that described in Chapters 3 and 4. A lookup table can be populated by actual measurements on calibrated samples or by accurate simulations of the forward problem. In the former case the advantage is that unknown systematic deviations of the experimental setup are automatically taken into account in the process. On the downside, the applicability of an experimental LUT remains restricted to a specific setup implementation, and its accuracy will be limited by that of the calibration of the phantoms used. Monte Carlo simulations, on the other hand, offer absolute control and flexibility on all the parameters, and can be used to explore large volumes of the parameter space which could be less straightforward to access experimentally. In the next sections, we describe the main features of a new Monte Carlo software (developed in close collaboration with Dr. G. Mazzamuto) focused on the spatio-temporal description of light transport, and how we used it to perform an in-depth characterization of the slab geometry case. While enabling the study of a wide range of parameters, its flexible implementation streamlined the building of an improved lookup-table for the retrieval of the optical properties of a turbid layer in the scattering regime where the diffusive approximation starts to fail.

5.2 MCPLUSPLUS: a scriptable Monte Carlo library for radiative transfer

All simulations performed in this work have been performed with a new Monte Carlo software library called MCPLUSPLUS that was developed from scratch to improve over existing Monte Carlo solutions. Indeed, several alternative implementations of the Monte Carlo method have been presented in the literature, each with different strengths and features. Notable examples include the early MCML [53] and its massively parallel reimplementations CUDAMCML and GPU-MCML [54, 184, 185] (on which MCPLUSPLUS is based), as well as other softwares modeling light polarization [186] or complex meshes [187]. As a qualitative validation of MCPLUSPLUS, a comparison of data simulated with MCPLUSPLUS and CUDAMCML is shown in Figure 5.1.

From a software design perspective, light transport exhibits two main features that must be considered. The first is represented by the structure of the building blocks of the problem, which lend themselves perfectly to be described as distinct *objects* with similar properties. For this reason, we developed our library in C++, an object-




Fig. 5.1.: Comparison of the results obtained for a slab with $L/l'_s = 10$ using MCPLUSPLUS and CUDAMCML. The original CUDAMCML code [184] has been modified to include the output relative to $w^2(t)$ shown in panel (c).

oriented programming (OOP) language which has already been suggested as an ideal programming language for this class of problems [188]. Since pieces of code can be encapsulated in reusable objects, OOP offers several advantages including scalability, modularity, ease of maintenance and abstraction. For example, taking advantage of the inheritance feature of OOP, creation of multiple classes with similar or overlapping tasks can be easily avoided, while their members and variables are easily shared, reducing redundant code. On the other hand, polymorphism allows to keep a high abstraction level referring to functions as generic entities rather than specific implementations (as in the case of different phase functions or spatial/temporal distributions). Equally important, is the natural propensity of OOP to describe a high-level *interface* to the software itself. Indeed, MCPLUSPLUS comes as a shared library rather than an executable package. A python interface to the library is also provided so that simulations can be easily set up and run through a simple script. Scriptability is one fundamental feature of MCPLUSPLUS as it easily allows to loop over the parameter space, which is particularly useful to automatize the building of a lookup-table or performing a fit. With respect to other scripting languages such as MATLAB, python is free and open-source software, yet offering a rich array of high-level libraries. Listing 5.1 shows an example of a python script used to set up and run a simulation such as that used to generate Figure 2.10. If desired, the raw output of the simulation can be organized in histograms where walkers emerging from the sample can be tallied up according to their point, angle or time of exit. Bivariate histogramming is also available, which is needed for example to access the spatio-temporal profiles. Walkers transmitted through the last boundary are classified either as TRANSMITTED or BALLISTIC if they have undergone at least one or no scattering events. Similarly, walkers leaving the sample from the first boundary are called REFLECTED or SPECREFLECTED. More than one class of walkers can be merged in the statistics using bitwise OR in the argument of `setWalkerTypeFlags`.

The sample is described as a sequence of layers of arbitrary materials. Exploiting ray-optics matrices, beam sources such as the `GaussianRayBundleSource` (see Figure 4.3) can be focused exactly to the desired depth even in the presence of intervening

Listing 5.1: Python script example setting and running the parallel simulation of 10^{10} walkers on 8 threads. As many output histograms can be defined and included in the output. Official units for angles, lengths and angles are degrees, μm and ps, respectively.

```

from pymcplusplus import *                                # import bindings to the C++ library

mat = Material()                                         # define scattering material
mat.n = 1.5                                              # set refractive index
mat.g = 0.6034                                           # set scattering anisotropy
mat.ls = 39.66                                           # set scattering mean free path

env = Material()                                         # define external material
env.n = 1                                                # external refractive index

sample = Sample()                                       # define sample instance
sample.addLayer(mat, 1000)                              # add 1000um-thick scattering layer
sample.setSurroundingEnvironment(env)                  # set external environment

source = PencilBeamSource()                             # define source type

sim = Simulation()                                       # define simulation object
sim.setSample(sample)                                   # add sample
sim.setSource(source)                                   # add source
sim.setNWalkers(10000000000)                            # set number of walkers
sim.setNThreads(8)                                       # set number of threads
sim.setSeed(0)                                           # set seed for the PNRG

hist = Histogram()                                       # define a histogram object
hist.setDataDomain(DATA_K, DATA_TIMES);               # set (bivariate) domain
hist.setWalkerTypeFlags(FLAG_TRANSMITTED);             # detect transmitted walkers
hist.setMax(90, 1000);                                  # set 90deg and 1000ps as upper boundaries
hist.setBinSize(0.5, 1);                                # set bin size for each domain
hist.setName("k_vs_times_transm");                     # set dataset label
sim.addHistogram(hist);                                  # append histogram to the simulation

sim.setOutputFileName("example.h5")                     # set binary output filename
sim.run()                                                # run simulation
sim.clear()                                              # free memory

```

layers. Several utility distributions are implemented, that can be used to draw angles, entrance positions and time delays according to the experimental configuration. The full hierarchy of classes and the definition of their methods is described in the online documentation [143]. The output of the simulation is saved in HDF5 format, a widely used format devised for storing large datasets in binary form.

The second relevant point for the design of a Monte Carlo software is given by the so-called ‘embarrassingly parallel’ nature of radiative transfer. All random walk trajectories can be computed independently of each other, meaning that any increase in the number of processing units contributes linearly to the efficiency of the simulation. For this reason, virtually all modern MC software take advantage of parallelization heavily, especially on GPUs where one can easily find thousands of cores. While offering multithreading capabilities, we developed MCPLUSPLUS to be run on CPUs rather than GPUs. In fact, while delivering the highest computing speed, graphical processing units present still few difficulties and limitations with respect to conventional programs [54, 185]. Conversely, the aim of MCPLUSPLUS was to deliver high precision, numerical accuracy, reliability and reproducibility. In fact, in a LUT approach, simulations and data

evaluation happen asynchronously, and therefore computing speed does not represent a major limitation. We must note however that, despite being CPU-based, MCPLUSPLUS still offers performances close to GPU-based solutions by exploiting multithreading on e.g., a multi-core workstation. Finally, CPU code also ensures complete hardware compatibility, while GPU software are hardware or even vendor specific.

A final relevant consideration is related to the magnitude of the simulations to be performed. The building of the LUT and the study of optically thin samples in particular, requires the simulation of extremely large number of energy packets. As we have seen in subsection 2.1.6, this involves the generation of a huge amount of random numbers, which must be done carefully. Within a computer, *pseudorandom* numbers are more efficiently produced using a pseudo-random number generator (PRNG), which is a deterministic algorithm that produces a deterministic sequence of numbers approximating the statistical properties of a truly random sequence. A PRNG can be started from an arbitrary initial condition using a *seed* state, and will always produce the same sequence while evolving its internal state from a given seed state. After a certain predetermined number of iterations (defined as its period) the algorithm will loop back to its initial state and start again to output the same sequence.

One trivial problem that arises with this approach is of course that the amount of random numbers drawn during the simulations should be kept way below the period of the PRNG to avoid useless repetition of identical trajectories. This is usually not an issue since state-of-the-art PRNGs such as the Mersenne Twister can have periods $> 10^{6000}$ at the cost of a moderate memory footprint. Still, this somehow prevents a widespread use of such high-level PRNGs on GPUs, where each thread has only a limited amount of memory available. However, this limitation is being rapidly mitigated by the advancement of both software and hardware platforms.

A much subtler issue, that to our knowledge has not been addressed yet in existing MC solutions, is connected to how computers represent real numbers in floating-point notation, which can heavily impact the generation of exponentially distributed numbers needed for the step length distribution (2.89). In a computer, any number is represented as a finite sequence of bits. Therefore, only a finite list of numbers can be represented and used in calculations. The main strategy to generate random numbers in accordance to a given distribution is to use its inverse cumulative distribution and feed it with random numbers $\xi \in [0, 1)$. Recalling the example of the exponential distribution, the useful relation for the step length is given by equation (2.61)

$$\ell = -l_s \ln(1 - \xi). \quad (5.1)$$

Due to the fact that only a finite number of representations $0 \leq \xi < 1$ can exist, the resulting exponential distribution is necessarily truncated at some $\ell_{\max} = -l_s \ln \varepsilon$, where we have assumed that $1 - \varepsilon$ is the closest number to 1 that can be represented. This truncation becomes statistically significant if one draws a large number of steps. In all modern CPU and GPU-based MC software implementations, the truncation length happens around 22.18 times the average step length (a thorough discussion of the

truncation problem is contained in Appendix B). This means that the probability of performing a step with $\ell > 22.18 \times l_s$ is identically null, while it should be

$$P(\ell > 22.18l_s) = \int_{22.18l_s}^{\infty} \exp(-\ell/l_s)/l_s d\ell = \int_{22.18}^{\infty} e^{-\ell} d\ell = \frac{1}{2^{32}}. \quad (5.2)$$

Consequently, whenever we run a simulation where more than $2^{32} \sim 4 \times 10^9$ steps need to be drawn, the standard implementation of the step-length distribution (SLD) will start to be appreciably biased towards shorter step lengths. It should be noted that, by definition, the number of steps drawn is always strictly bigger than the number of simulated energy packets — often by few orders of magnitude — so that this issue is not as remote as it might seem. Nonetheless, for most applications, a smaller number of energy packets is usually sufficient and this effect can be almost neglected. Driven by the more fundamental nature of our investigation, we developed MCPLUSPLUS providing it with the possibility to use an exponential distribution where the truncation probability is pushed below 2^{-64} , meaning that more than $\sim 5 \times 10^{20}$ steps can be generated without incurring in a biased step-length distribution. As described in detail in Appendix B, this requires the use of a 64-bit PRNG in combination with a 128-bit long double floating point representation for the uniform random variate ξ . The implementation of these features ultimately determined our choice to develop MCPLUSPLUS on the CPU, where they can be straightforwardly implemented. In contrast, forcing GPU software to meet such massive scale and precision requirements would be extremely unfavorable, given its specific vocation for lower-precision arithmetic.

5.3 Deconstructing light transport at the ballistic-to-diffusive transition

As we discussed in the introduction, light is a particularly versatile tool to investigate the optical properties of materials in many different fields of science, medicine and technology. Each optical parameter provides different insights on the sample composition, such as the density and scattering strength of its constituents, the concentration of pigments or absorbing chromophores, and the shape and size of the single scatterers. Accurate retrieval procedures for these parameters are therefore of paramount importance in a number of applications.

When considering the slab geometry, the diffusion approximation (DA) represents a straightforward and widely used theoretical framework, in that it provides a complete set of analytical expressions to describe light transport both in the spatial and temporal domain. However, as we have shown experimentally in Chapter 4 for a variety of samples, its validity range is limited to highly turbid materials and DA-based inversion techniques are affected by systematic errors especially when applied in the regime where the optical thickness (OT) falls below ~ 10 . For this reason, despite the fact that intermediate-turbidity layers are of large interest both in fundamental research [68, 84,

189, 190] and applications [91, 191–194], this class of samples remains to date less studied given the need for non-approximated numerical techniques. Additionally, even when dealing with turbid materials, the samples may exist only in a small range of thicknesses and sizes, which is often the case in the biomedical field when studying tissues as the ocular *fundus* [91], vascular walls [195], cellular cultures [196], skin dermis [197], skull bones [198] or dental enamel [199] — to name a few.

In all these circumstances, experimental data evaluation must rely on more accurate methods such as refined approximations of the radiative transfer equation (RTE) or Monte Carlo (MC) simulations. Nonetheless, due to its simplicity, the diffusion approximation still retains a large appeal, and great efforts are constantly made to extend its validity range introducing all sorts of *ad-hoc* modifications [70, 74, 77]. Still, an intriguing conundrum exists when applying the diffusive picture to optically thin media. On one hand, in low-turbidity samples the diffusion approximation is bound to breakdown in the sense that light transport will be dominated by unscattered (ballistic) light, or light undergoing too few scattering events. On the other hand, however, if a time-domain analysis is available, early light can be rejected focusing solely on the fraction of light that propagated deeply in the multiple scattering regime even in samples which would not be typically associated with this transport regime [194, 200]. We would therefore expect that this late-time component fulfills the validity assumptions of the DA even in the thin slab geometry.

In this section, we elucidate this point by performing an extensive Monte Carlo characterization of light transport over a wide range of optical and geometric parameters, with the aim of testing the validity of the diffusion approximation for decreasing optical thickness. In particular, as we have seen in Chapter 2, diffusion theory casts an incredibly simple prediction on *transverse* transport which could be profitably applied in a thin slab geometry (equation (2.104)), given that boundary and confinement effects are less relevant along the slab's main extension. As discussed in Chapter 4, strong numerical and experimental evidence exist that studying transport along transverse directions enables a description that is more robust against experimental uncertainties and easier to interpret within the diffusion approximation [110, 157, 158, 201], as opposed to other methods based on the spatially integrated transmission or its decay constant τ . Despite this, the vast literature available for the slab geometry typically focuses on axial rather than transverse transport [68–70, 72, 84, 85, 102, 189], supposedly because the latter was difficult to access experimentally in a direct way up to very recently. In this respect, our investigation fills the gap with the extensive characterization already available in the literature regarding the breakdown of the diffusion model along the axial direction, finally providing the richer 3-dimensional picture of the transition between the ballistic and the diffusive regime. This characterization allows to identify a set of transport descriptors enabling the building of a new, robust type of lookup table whose advantages and working principles will be presented in the next Section.

The typical configuration that we studied is that shown in Figure 5.2, with a pencil beam pulse $\delta(\mathbf{r})\delta(\mathbf{k} - \mathbf{k}_i)\delta(t)$ with $\mathbf{k}_i = (0, 0, 1)$, impinging normally on a infinitely extended scattering slab. According to the DA, this typically results in a Gaussian transmission profile with a standard deviation growing linearly as $w^2(t) = 4D_{\text{DA}}t$ with




Fig. 5.2.: (a) Sketch of the investigated configuration. An infinite slab is illuminated by a pencil beam pulse. Transmitted light is collected at different times and positions. A few trajectories and normalized transmitted intensities are presented for illustrative purposes in the case of an optically thin, index matched slab $L = l'_s = 1$ mm (i.e. optical thickness = 1), shown to scale. The scattering anisotropy g is set to 0. (b) An approximately Gaussian profile is transmitted at each time slice, whose mean square width and integrated intensity respectively grows linearly and falls exponentially.

a slope determined by the diffusion coefficient $D_{\text{DA}} = l'_s v / 3$. This quantity is more generally referred to as the mean square width (MSW), which can be defined for an arbitrary intensity distribution $I(\rho, t)$ through the relation (2.103)

$$w^2(t) = \frac{\int_0^\infty \rho^2 I(\rho, t) \rho \, d\rho}{\int_0^\infty I(\rho, t) \rho \, d\rho}. \quad (5.3)$$

Analogously, the integrated transmittance exhibits a typical exponential decay with a time constant given by the relation (2.109)

$$\frac{1}{\tau_{\text{DA}}} = \frac{\pi^2 D_{\text{DA}}}{(L_0 + 2z_e)^2} + \mu_a v \quad (5.4)$$

which depends explicitly on the physical thickness, the absorption coefficient, and the boundary conditions through the extrapolated length z_e .

In the simulations, energy packets propagate inside the scattering medium following a standard random-walk algorithm for multi-layered structures implemented in MCPLUS-PLUS. Scattering step lengths are drawn following an exponential distribution while scattering angles are generated using the well-known Heyney-Greenstein function (2.30). For each transmitted and reflected packet, the time and point of exit from the sample is recorded. Assuming a fixed slab thickness of $L_0 = 1$ mm we simulated different values of the optical thickness (OT) ranging from 1 to 10 by varying the reduced scattering mean free path l'_s between 0.1 mm and 1 mm. In this case, the reciprocal of the optical thickness OT^{-1} is considered as the main parameter, representing a dimensionless figure

for the reduced scattering mean free path l'_s . Each fixed-OT simulation has been run for 11 different values of the anisotropy factor g between 0 and 0.99 and 16 values of the refractive index contrast n ranging from 0.6 to 2.2, for a grand total of 2816 simulations of 10^9 photons each. The full scriptability of MCPLUSPLUS largely streamlined the accumulation and sorting of the simulations. For the sake of convenience, since Fresnel reflection coefficients depend solely on the relative refractive index contrast $n = n_{\text{in}}/n_{\text{out}}$, we kept $n_{\text{in}} = 1$ constant while varying n_{out} in order to have a consistent time scale over the whole set of simulations. The real time scale of any single simulation can be recovered by simple multiplication by the actual value of n_{in} .

The output of the simulations is analyzed in terms of the transmittance MSW growth rate and the total transmittance decay time, as summarized in Figures 5.3a and 5.3b, respectively. It should be noted that these descriptors converge exactly to the same respective values also when considering reflectance, and are therefore valid for both types of measurements. In order to build the hyper-surfaces in the parameter space, we start by evaluating the exponential decay time of simulated data in order to set, for each simulation, a time-scale normalized to τ . Both the decay constant and the MSW rate are retrieved by fitting the data only in a temporal range between $4\tau < t < 9\tau$, which ensures that the fitting is performed at times long enough to extract the asymptotic values of the parameters. Moreover, this adds consistency to the fitting method itself between different simulations. In the case of the MSW slope, the limited fitting range allows to systematically exclude early-time light transmitted before the onset of the diffusive regime, while the upper limit helps avoiding the noise found at very long times due to insufficient statistics. It might be called into question whether it is appropriate to use the decay rate τ as a time unit for the MSW evolution, since the former is mainly determined by transport properties along the thickness direction, while the latter occurs along the plane. A decay-time-based temporal range provides indeed a convenient way of defining a consistent, self-tuning fitting window across the whole dataset. This simple choice is also advocated under practical reasons, since the decay time is undoubtedly the actual temporal unit that eventually dictates — both in real and numerical experiments — the signal-to-noise ratio. In this respect, every diffusion coefficient within our simulated phase space has been determined under equal noise conditions. No less important, limiting our investigation to a long-time window is also relevant under a more technical point of view: i.e. it renders irrelevant for all practical purposes the specific choice of both the spatial source distribution and the phase function.

Values of τ and D obtained for each simulation are eventually arranged in the form of a hyper-surface as shown in figures 5.3a and 5.3b, respectively. In order to neutralize the noise originating from statistic fluctuations and fitting uncertainty, we consider each simulated n -slice separately and smooth the data through a Loess fitting routine (range parameter set to 0.25) as shown in the third row of plots. Smoothed slices are eventually reassembled together to perform a cubic interpolation along the index-contrast axis to obtain a hyper-surface for D and τ that can be evaluated seamlessly for any triplet in the (n, g, OT^{-1}) parameter space. Interpolation has been performed separately on the $n \leq 1$ and $n \geq 1$ regions of the parameter space due to the sharp first-derivative discontinuity occurring at $n = 1$.

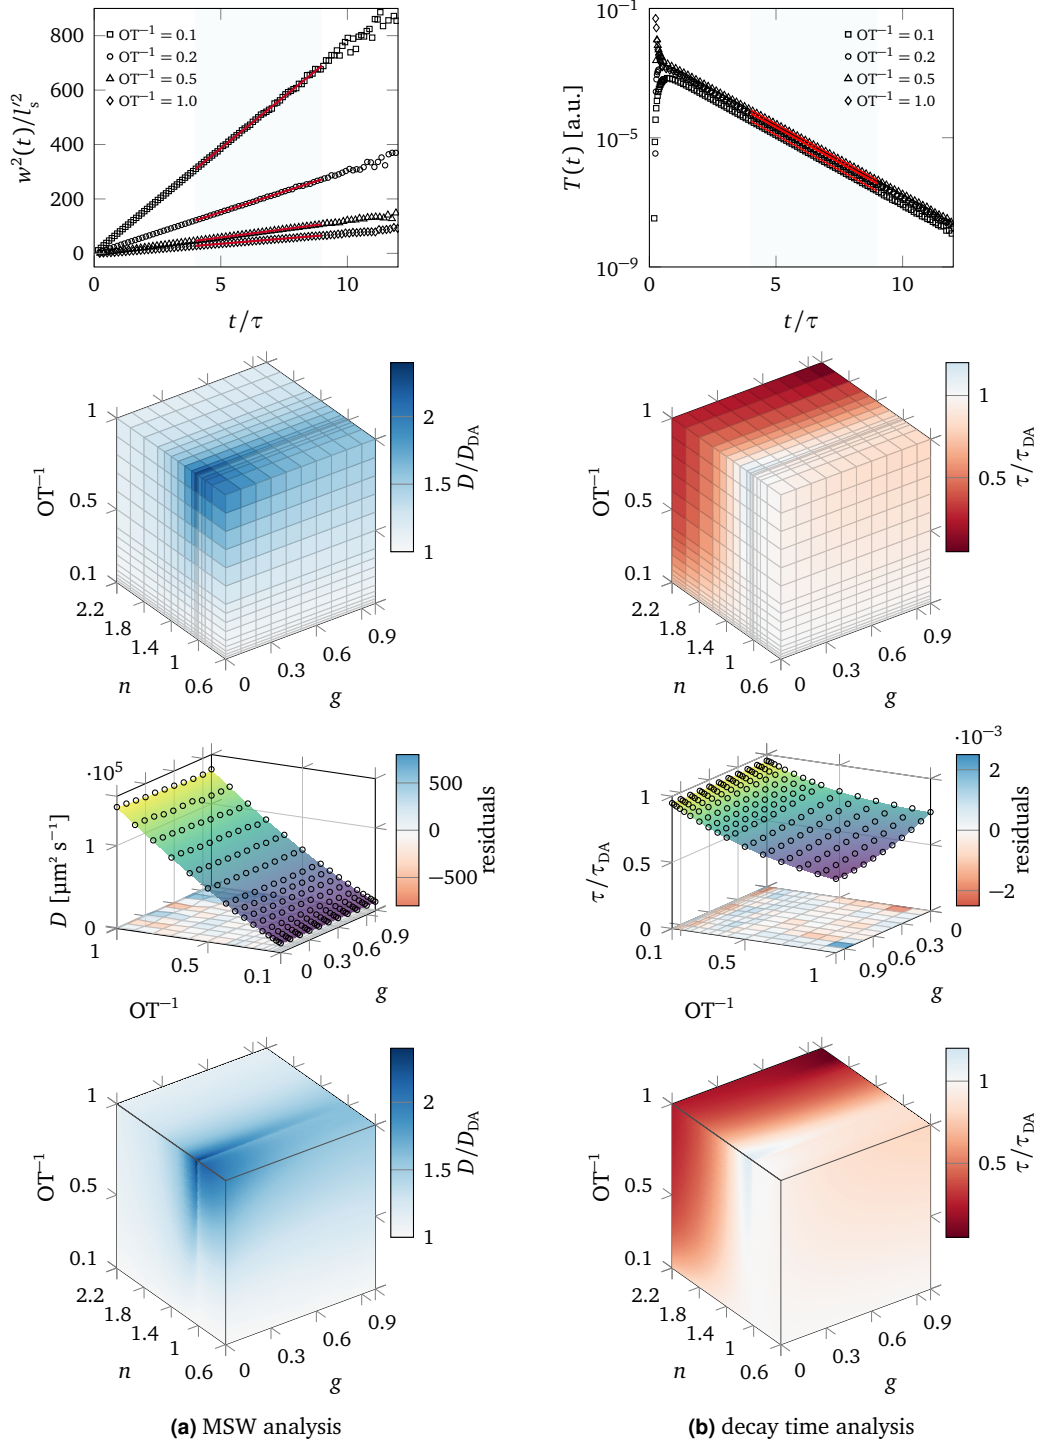


Fig. 5.3.: From top to bottom, steps followed to generate the hyper-surface of (a) relative MSW slope and (b) decay-time deviations. Upper panels: subset of simulated time-resolved MSW and total transmittance curves for $n = 1.4$ and $g = 0.9$ at different values of l'_s/L_0 . Curves are fitted over a temporal window ranging from 4 to 9 decay times. Discretized hyper-surfaces are obtained with the relative deviations of the investigated parameters with respect to the DA. Each simulated n -slice ($n = 1.4$ shown) of the parameter space is processed through a Loess fitting routine. The smoothed slices are reassembled together to carry a gridded interpolation along the index contrast axis.

5.3.1 Mean square width expansion

The upper panel of Figure 5.3a shows a subset of simulated MSW data obtained for typical optical properties of interest in the bio-optical field ($n = 1.4$, $g = 0.9$), which surprisingly exhibit a linear asymptotic increase even at the lowest value of the optical thickness. As previously discussed, the value of the MSW at each instant is exactly independent of absorption, which has been therefore ignored from the simulations. The retrieved values of D , evaluated as $1/4$ of the variance slope, have been normalized by the expected value $D_{DA} = l'_s c / 3$ and arranged in a hyper-surface of relative deviations. The obtained volume is sampled in a discrete set of points in the (n, g, OT^{-1}) space, with a 1 : 1 correspondence with the number of performed simulations. Noise coming either from limited statistics or fitting uncertainties is largely suppressed by applying a local regression algorithm using weighted linear least squares and a 2nd degree polynomial model as provided by the Loess MATLAB model (Figure 5.3, third row). This allows to obtain an accurate, seamlessly sampled volume suitable for finer interpolation, as shown in the final step of Figure 5.3a.

A few comments are in order. Firstly, the present investigation is intended to focus on long-time/asymptotic transport. To this purpose, the diffusion coefficient D has been evaluated by the linear slope of the mean square width (MSW) in a time window ranging from 4 to 9 decay times, as determined from time-resolved curves. Depending mainly on the optical thickness of the sample, there is an early-time range where the MSW exhibits a super-linear increase. We carefully checked that the aforementioned fitting time range was always largely excluding such non-linear time range in order to address safely the asymptotic slope.

Secondly, it is well known that most biological soft tissues share a refractive index equal or close to $n_{in} = 1.4$ [202]. This is supposedly the reason why refractive index variations have so far been disregarded in similar multi-parameter investigations [93, 167, 178, 182]. Nonetheless, we included the refractive index contrast as a simulation parameter because, especially in the case of thin slabs, the range of interest for n is undoubtedly wider. The case of small n can be of interest for samples that are enclosed in glass slides, or laid or immersed in different substrates/solutions, whereas the higher values have been included envisioning possible applications of our study to metal oxides and similar highly scattering materials, which are extremely relevant, for instance, for coatings and in photovoltaics.

Looking at the obtained data, two features are immediately noticeable. Firstly, the diffusion approximation appears to always *underestimate* the actual spreading rate, of course recovering agreement for higher optical thicknesses as expected. A second, finer feature occurs in the close proximity of $n = 1$, particularly evident at low g and OT values. Both these features arise from the interplay between geometric and boundary conditions. In particular, the presence of internal reflections in a thin layer geometry helps to selectively hold inside the slab those energy packets that happen to draw statistically longer steps, as we discuss in more detail in the following Chapter. To the purpose of solving the inverse problem, it is worth noting that the MSW slope exhibits a distinct pattern of characteristic deviations from the DA, which can be therefore exploited

as a guide to unambiguously retrieve the *intrinsic* microscopic transport properties of a given sample.

5.3.2 Decay time and absorption

The upper panels of Figure 5.3b show respectively a typical set of time-resolved transmittance decays and the hyper-surface of relative deviations from the DA predictions. Two main features are worth commenting when comparing these results to the previous MSW characterization. First of all, the obtained decay time deviations are more significant, reaching down to just a few percent of the expected value for the highest contrast n and anisotropy factor g . It is indeed known that $n > 1$ refractive-index contrasts are more difficult to be taken into account even when applying appropriate boundary corrections and even at higher optical thicknesses (see, for example, Figure 2.10). Secondly, as opposed to the MSW case, deviations in both directions are possible, with the τ/τ_{DA} ratio taking values both above and below 1. This might help explaining some experimental evidences obtained in thin disordered samples that are still debated at a fundamental level [71, 85, 87, 203–205], as we will further discuss in the next Chapter. These findings stress the importance of an accurate and precise modeling of the index contrast, which we think has been often overlooked, for example when a symmetric averaged contrast is used to model asymmetric experimental configurations [71, 85, 204].

Despite the vast literature regarding the validity range of the diffusion approximation in the time domain [67–69, 72, 102, 147, 150, 206], a comprehensive understanding of the interplay between optical thickness, refractive-index contrast and absorption is still object of debate. It is a common assumption that the diffusion approximation fails gradually with decreasing optical thickness, with $\text{OT} = 8$ being customarily considered as the lower threshold under which the introduced error starts to be significant [72]. Nevertheless, as we have seen in Section 4.2, even in the absence of absorption a $\text{OT} > 8$ slab sample with $n \sim 1.5$ can exhibit a transmittance decay time such that the diffusion approximation is unable to provide *any* real solution at all (see Figure 4.4c), thus suggesting that the breakdown of the diffusion approximation might step in abruptly depending on the interplay between different parameters other than the optical thickness. As we will show in the next Section, as expected, the experimentally observed deviations are in perfect agreement with our new set of simulations.

A few words should be spent on the role of absorption, which we have excluded from the simulations even though it contributes to defining the value of the transmittance decay time through equation (5.4). This does not imply any loss of generality, since also in the case of the decay time, the presence of absorption can be accounted for *exactly* by shifting

$$\frac{1}{\tau_{\text{DA}}} \rightarrow \frac{1}{\tau_{\text{DA}}} + \mu_a \nu, \quad (5.5)$$

τ_{DA} being the decay time in the non-absorbing case.

The problem with absorption is that both scattering and absorption can deplete specific intensity from a given position, time and direction (an effect sometimes referred to as absorption-to-scattering cross-talk). Hence, retrieving its unknown value from

experimental data has been to date a challenging task. In a quest for decoupling their effects, crude approximations have been introduced even in the time domain, relying on the assumption that the intensity decay time would eventually become independent of the scattering coefficient at long time scales [207, 208], which however can lead to wrong estimations unless an extremely large dynamic range is available [176, 209].

Besides that, it is often reported that the diffusion approximation is expected to hold only for weakly absorbing media since the onset of the properly diffusive regime requires long trajectories to contribute dominantly to transport properties, whereas these are selectively suppressed by absorption [60]. This explains why absorption is often considered as a major hindrance in the correct assessment of transport properties [79, 210, 211], if not even an invalidating condition for certain optical parameter measurements [38, 212, 213]. For this reason, techniques capable of directly accessing the MSW recently aroused a great deal of interest given the absorption-independent nature of the variance expansion [42, 110, 148, 149, 157] that allows for the first time to decouple exactly absorption from scattering. The full potential of MSW measuring techniques is still to be fully unraveled: as we will demonstrate in the following, it can play a key role in accurately retrieving both parameters.

5.4 A Monte Carlo LUT based on spatio-temporal descriptors

As we have seen, the solutions to the RTE in the range of optical thicknesses comprised between 0 and 1 exhibits some significant deviations from the DA and the similarity relation, while at the same time retaining its main hallmarks such as the steadily linear MSW growth rate. This results in a pair of multidimensional hyper-patterns, the joint evaluation of which represents a characteristic signature of a given set of optical parameters. This observation leads naturally to the definition of a lookup-table approach.

The main feature of the LUT routine that we designed is that, for the first time to our knowledge, it relies on observable quantities that do not require any absolute measurement and are well into the multiple scattering regime. This offers several advantages over existing solutions mentioned in Section 5.1.

- both the asymptotic decay time and the mean square width slope can be measured without any reference to the excitation intensity, therefore there is no need for an absolute calibration of the source or the detector. Decay time determination is also not strictly connected to any particular detection geometry, which stands in contrast with other typical techniques often requiring a particular configuration of collection fibers, integrating spheres or angular measurements.
- because of the asymptotic nature of both τ and $w^2(t)$, the actual temporal response function or spatial excitation spot size are eventually irrelevant to their accurate determination
- precise determination of the origin of the time axis (i.e., the exact time of pulse injection), while being dramatically relevant in many analogue situations (see

Figure 4.2), is here made completely irrelevant since both the decay time and the linear increase of the mean square width do not exhibit any critical dependence on the exact delay at which they are determined, provided that it is sufficiently large.

- with respect to MC-based fitting routines, a lookup-table routine is more suitable for real-time solving of the inverse problem since it does not involve any iterative procedure. While this guarantees ideal performance, on the downside we must note that it is less clear how to define the uncertainty of retrieved values. Typical approaches involve mapping the relative error on the retrieved parameters over a broad range of independent simulations, in order to give a numerical estimation.
- several issues typically associated to fitting routines are also obviated. Once that the two scalar descriptors are calculated with the proper, original binning, they can be virtually rescaled arbitrarily without the risk of introducing any binning-related artifact.
- the problem of correct bin positioning is also removed. Midpoint positioning adopted in our case represents an exact solution for the linear increase of the MSW. While this is not the case for a monoexponential decay, it can be again shown trivially that midpoint positioning does leave the decay constant exactly unmodified.
- as we will show in the following, a possible use of our LUT routine is that of retrieving μ_a and l'_s . It must be nonetheless stressed again that none of the simulations that compose our LUT need to include the effect of absorption, which stands in contrast with the customary practice of most LUT approaches demonstrated to date. This allows to deal with a simulation phase space of reduced dimensionality without any loss of generality, which represents a enormous saving on the computational burden of MC simulations.
- finally, since there is no need to directly simulate absorption *nor add it after the simulation*, it is also not necessary to store exit times and positions on a single-packet basis. This allows to hugely reduce the output size for each simulation and streamline its handling, thus allowing for larger statistics to be collected.

At this stage, our demonstrative MC-LUT routine is based on two descriptors and therefore allows to retrieve only pairs of transport parameters, e.g., l'_s and g (and consequently l_s) assuming that absorption is known, or l'_s and μ_a assuming a known value of g (which is a common practice in similar works, especially those involving biological samples [93, 166, 167]). The effective refractive index and the thickness of the layer are also expected as input parameters. A freely queryable version of the full dataset is available online with a dedicated interface [214]. To illustrate the steps involved in the LUT routine, we test the retrieval procedure on two simulated samples with $L = 1.3$ mm, $n = 1.38$, $g = 0.95$, $l_s = 45$ μm and μ_a respectively equal to 0.2 mm^{-1} and 0 mm^{-1} (Figure 5.4). From these simulations we extract the mean square width slope and the decay time, which are used as inputs for the LUT routine.

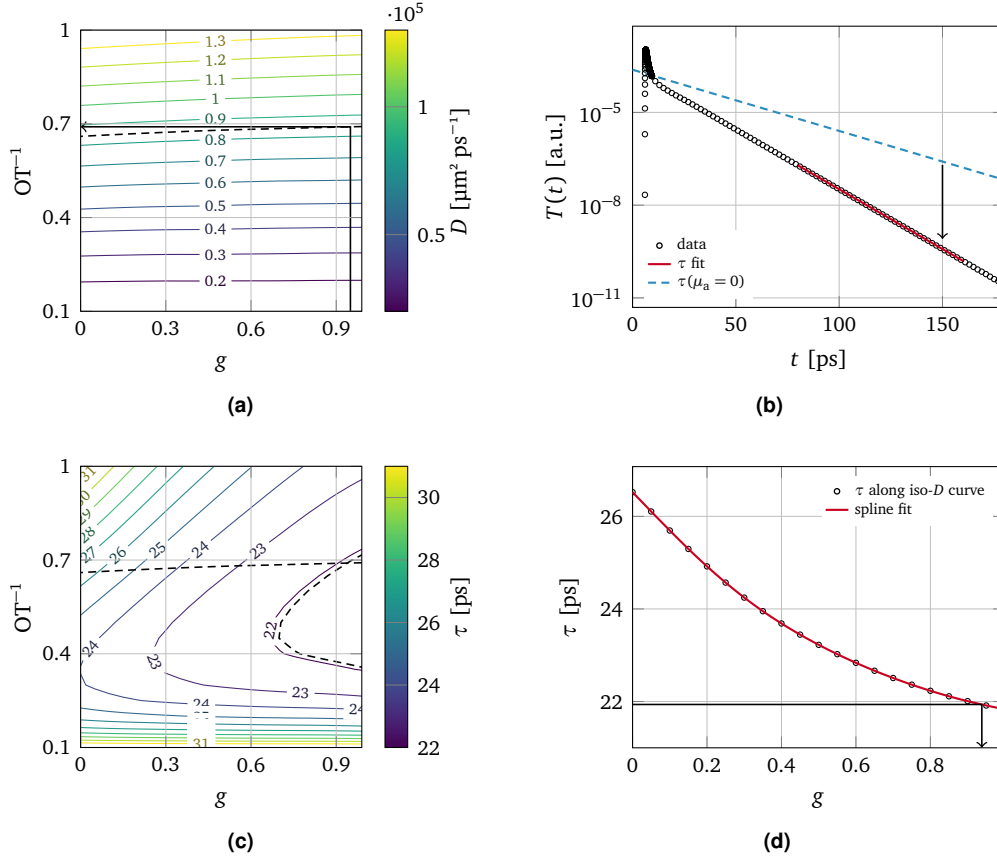


Fig. 5.4.: Demonstration of the MC-LUT routine. (a)-(b) Retrieving l'_s and μ_a with known g . The value of D retrieved from the MSW slope is marked on the corresponding n -slice of the parameter space (dashed line). The intersection with the known value of g provides the best estimate for OT^{-1} . The expected decay time is read in the LUT at position $\tau(n, g, OT^{-1})$ and compared to the experimental one to retrieve μ_a . (c)-(d) Retrieving l'_s and g with known μ_a . Intersecting the iso- D and iso- τ curves yields the estimated value of g .

As a first step of the LUT procedure the MSW and decay time hyper-surfaces are (exactly) rescaled both in time and space to match the target thickness and refractive index of the investigated sample. The original simulations were performed for a sample of thickness $L_0 = 1$ mm and unitary internal refractive index; dimensional analysis shows that eventually the mean square width and decay-time hyper-surfaces are to be rescaled by $L/(L_0 n_{in})$ and $n_{in} L/L_0$ respectively. Successively, the interpolated hyper-surfaces are sliced at the known value of the refractive-index contrast.

Let us start by considering the case where we assume a known value for g and try to retrieve l'_s and μ_a . A linear fit of the MSW data returns a growing rate of $337\,750 \mu m^2 ps^{-1}$, corresponding to a $D = 84\,437 \mu m^2 ps^{-1}$. This value can be represented as an iso- D level on the sliced (g, OT^{-1}) -surface, which yields directly the value of OT^{-1} by intersection with $g = 0.95$ (Figure 5.4a). Notably, in case the scattering anisotropy is not known in advance, plugging into the LUT reasonably bounded values helps getting an estimate of how an uncertainty on g spreads over l'_s and eventually μ_a .

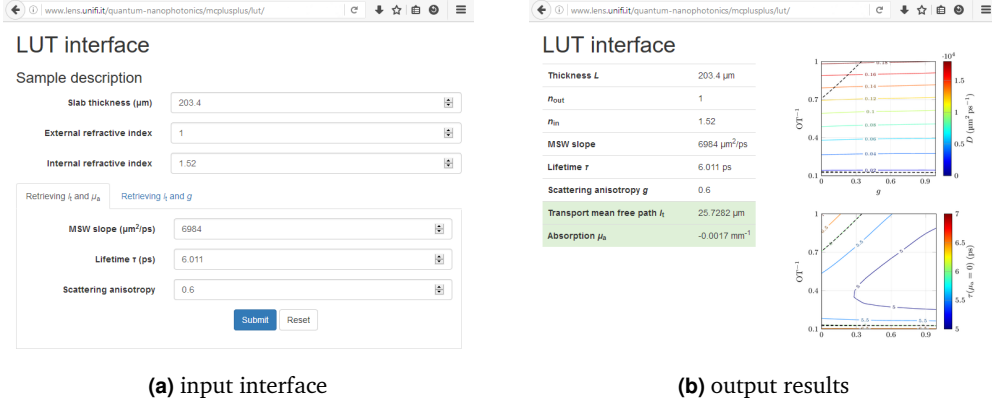


Fig. 5.5.: Online LUT input and output interfaces, available at <http://www.lens.unifi.it/quantum-nanophotonics/mcplusplus/lut/>. Input data are taken from the homogeneous scattering sample of Figure 4.1.

Following OT^{-1} determination, it is sufficient to read the expected absorption-free decay time stored in $\tau(n, g, OT^{-1})$ from the interpolated decay-time surface and compare it directly to the experimental value. The discrepancy between their reciprocal values will directly give $\mu_a v$ through equation (5.5) (Figure 5.4b). Fitting the simulated transmitted intensity decay yields a decay time of 11.234 ps from which we finally retrieve $\mu_a = 0.1997 \text{ mm}^{-1}$ and $l'_s = 897 \mu\text{m}$, to be compared with the nominal values of $\mu_a = 0.2 \text{ mm}^{-1}$ ($\delta \approx -1.5 \times 10^{-3}$) and $l'_s = l_s / (1 - g) = 900 \mu\text{m}$ ($\delta = -0.3 \times 10^{-3}$).

The second implementation of the routine allows to retrieve l'_s and g assuming that μ_a is known. A common case is that of vanishing absorption, as it was the case for example with the samples studied in Chapter 4. The evaluation of D is unaffected and yields the same result of the absorbing sample. Now the corresponding iso- D curve can be superimposed on the $\tau(n = 1.38)$ surface along with the experimental non-absorbing decay time of 21.936 ps (dashed line in Figure 5.4c). Their intersection, which can be calculated for example by spline interpolation along the iso- D line as shown in Figure 5.4d, finally gives the estimated value of g . In this case we obtained $l'_s = 897 \mu\text{m}$ and $g = 0.938$ ($\delta = -1.2 \times 10^{-2}$).

It is worth testing our LUT against the experimental data relative to the homogeneous scattering slab that we characterized in Chapter 4 measuring both its MSW expansion rate and the total transmittance decay time. Figure 5.5 shows the output of a LUT query performed on the online interface, assuming a value of $g \sim 0.6$ for the TiO_2 nanoparticles. We obtain an estimate of $l'_s = 25.7 \mu\text{m}$ for the reduced scattering mean free path, in good agreement with the previously determined value of $25.5 \mu\text{m}$ retrieved with a brute-force fit. As regards the value of absorption, the output value is slightly negative, corresponding to a *gain* rather than an absorption length of few tens of cm. We interpret this result as a numerical fluctuation consistent with null absorption.

Thorough evaluation of errors should be performed on a wide range of parameters, both from simulations and experimental data, which is beyond the scope of this proof-of-concept demonstration. Nonetheless it is clear that, especially at lower thicknesses

where the diffusion approximation is more defective, our routine offers very accurate inversion capabilities as compared to other slab-geometry fitting and/or LUT approaches [167]. Uncertainties as low as a few percent with respect to simulated data have been demonstrated in other works for the semi-infinite geometry using integrated intensities as the input parameters. It might be questioned whether this kind of uncertainty is meaningful, since integrating-sphere measurements themselves suffer of both random and systematic errors of similar magnitude in the first place [93]. On the contrary, the slope of the mean square width and the transmittance decay time can be typically determined with better precision, accuracy and robustness, since their scalar value is a collective property of data points in a curve rather than the straight output of a single measurement.

As a last point, it is interesting to discuss possible extensions of our routine applicability. At least a third input descriptor in addition to the decay time and the MSW slope needs to be identified in order to retrieve simultaneously all three transport parameters at once from an unknown medium. A possible candidate could be represented by the asymptotic-tail decay of a steady-state transmission profile, which is also easily measurable and should exhibit a small but appreciable dependence on g at low optical thickness. For all practical purposes, this asymptotic decay rate would feature all the previously listed advantages, with the possible exception of the last one, because of the need to add absorption *ex post*. Other relative parameters could be exploited, taking advantage of their g dependence, such as the transmittance rising time [102], and many more if also diffusely reflected light is included into the analysis, given that it is more affected by low-order scattering events occurring near the source. Finally, the domain of the lookup-table could be easily extended to negative values of g which, albeit less commonly found in typical applications, are known to be possible even in random, uncorrelated assemblies of semiconductor scatterers [215].

To sum up, lookup-table methods are very general in their nature and consequently can be profitably applied in a number of practical use cases. Of course, in order to tackle more complex samples (e.g., multilayered or anisotropic slabs) more observables are needed. Nonetheless, we believe that, whenever possible, mean-square width and decay-time measurements should always be preferred and included in every LUT-based retrieval routine because of their intrinsic robustness.

Asymptotic transport in bounded media

6

” *I have yet to see any problem, however complicated, which, when you looked at it in the right way, did not become still more complicated.*

— P. Anderson
(writer)

6.1 Diffusive light transport in a semitransparent slab

During World War II, Abraham Wald (1902–1950) was working as a member of the Statistical Research Group at Columbia University when he was asked to estimate the vulnerability of military aircraft, so that reinforcement strategies could be devised to minimize losses [216, 217]. Previous studies from the Center for Naval Analyses, based on the examination of survived planes, recommended to further reinforce the parts of the aircraft that were damaged the most, which, however, proved completely ineffective, if not detrimental. Wald noted the fundamental flaw of the analysis: the aircraft population considered was strongly biased, since it relied solely on the analysis of survived aircraft, while there was no means to assess the damage of planes that had been taken down. The collected data was nonetheless providing valuable information. In fact, Wald proposed that protections should instead be added to those parts of the aircraft that were *not* hit, since the fact itself that the plane was able to return meant that the damages were non-critical.

This anecdote helps introducing the statistical analysis that we will discuss in this chapter. As in the case of Wald’s aircraft, light propagation in a bounded medium can also be considered as a ‘survival’ game where energy packets are terminated either by absorption or transmission at boundaries. Indeed, light that is held for longer inside a sample carries more and more accurate information about the average microscopic properties of the material. Nonetheless, it is interesting to check whether it exhibits any further statistical ‘signature’, akin to the bullet holes in Wald’s survived aircraft wings.

In the previous Chapters, we have discussed the role of time-domain measurements under several perspectives. Generally, a large time scale is always considered as the main figure of merit to evaluate the onset of the diffusive regime. However, it is rarely pointed out that reaching the diffusive regime does not directly imply that the diffusion approximation should hold better. In fact, we have seen that important deviations exist, and that they can be exploited to accurately characterize the transport properties of a

medium. In the following, we further explore this significant distinction and elucidate how it arises, as well as its effects and potential impact on the way we usually link microscopic scattering properties and their macroscopic, measurable counterpart. This is particularly true for the slab geometry that we analyzed so far. In fact, while on one hand it is common knowledge that diffusion theory cannot describe light propagation in thin layers, on the other hand even in an optically thin slab late-time transport will be eventually determined by a multiple-scattering process whose characteristics are still largely unexplored. Indeed, as we have seen in the previous Chapter, after a short transient, propagation along the slab plane converges towards a diffusive regime even at an optical thickness of 1, exhibiting a linear MSW growth. Still, when compared to the diffusive prediction, a significant deviation from the expected value is found.

To our knowledge, the validity range of the simple linear prediction $w^2(t) = 4Dt = 4\nu l'_s t/3$ has never been studied to date. The numerical results obtained in Figures 5.3a and 5.3b show that a peculiar pattern of deviations is concentrated in the proximity of $n = 1$. Figure 6.1a shows a representative series of curves taken on the upper surface $(n, g, 1)$ of the studied parameter volume, i.e., for a sample with $L_0 = l'_s = 1$ mm. The first remarkable feature that leaps out is that the fitted MSW slope is always greater than the value expected from diffusion theory, meaning that the diffusion coefficient appears to be enhanced with decreasing optical thickness. A first, qualitative explanation for this enhancement can be attempted based on the d -dimensional modeling of diffusion as a random-walk process (2.88), which, given a step-length distribution $p(\ell)$ with finite moments $\langle \ell \rangle$ and $\langle \ell^2 \rangle$, predicts a mean-square d -dimensional displacement growing as $2dDt$ with

$$D = \frac{1}{2d} \nu \frac{\langle \ell^2 \rangle}{\langle \ell \rangle} = \frac{1}{d} \nu l'_s \quad (6.1)$$

where the last equality holds for an exponential step-length distribution (SLD) with average step length l'_s [218]. As the optical thickness of the simulated slab decreases, transport occurs in an increasingly planar geometry. Hence, as suggested by equation (6.1), the effective diffusion coefficient D as inferred from the MSW slope might be up to $3/2$ times higher than its bulk nominal value. The perceived spatial dimensionality is also affected by the refractive index contrast. Near $n = 1$, any energy packet leaving the sample at long times will have performed an almost planar trajectory, akin to a purely two-dimensional walk. On the contrary, strong boundary reflections allow trajectories to fold back into the sample, which is therefore perceived more as a three-dimensional environment (which also explains why the diffusive approximation recovers gradually for high values of the refractive index contrast, see the middle panel of Figure 6.1a). A closer look at the data, however, shows that diffusion exhibits a local *minimum* at $n = 1$, rather than a maximum, with the D/D_{DA} ratio exhibiting a sharp modulation across the index-matching condition. Around $n = 1$, diffusion appears to be asymmetrically enhanced, reaching an absolute maximum around $n = 1.016$ for $g = 0$.

Interestingly, a similar behavior to what we described for the MSW is also found in the relative deviations of decay times from the diffusive prediction, as shown in Figure 6.1b, and is therefore not strictly limited to the propagation of light along the slab plane. This point deserves particular attention, especially given that decay time measurements

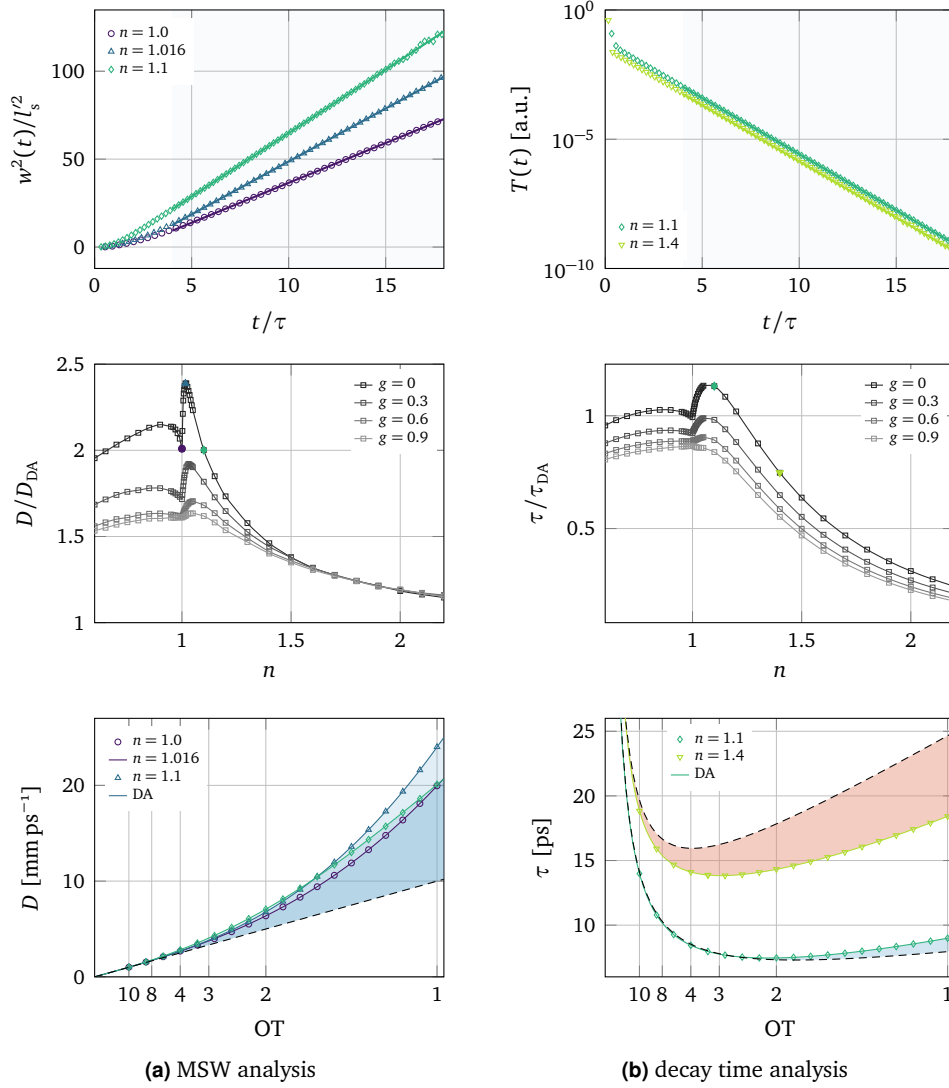


Fig. 6.1.: From top to bottom: MSW expansion rate and decay time for a slab configuration with $l'_s = L_0$, $g = 0$ and different refractive index contrasts, exhibiting (a) a perfectly linear growth and (b) a monoexponential decay after a short transient. Dependence on n and g of the diffusion coefficient D and decay constant τ as inferred from a linear fit of $w^2(t \geq 4\tau)$ and $\log T(t \geq 4\tau)$ for different samples with OT = 1. Solid points represent the values retrieved from the fits shown in the upper panel. Absolute positive (blue) and negative (red) deviations with respect to the DA prediction are shown in the lower panels for different optical thicknesses. Solid lines serve as guides to the eye.

of integrated transmittance have long been experimentally accessible and exploited to estimate the diffusion coefficient via equation (5.4). A similar dependence on n with respect to the previous case can be appreciated plotting the ratio between the decay time τ as fitted from the Monte Carlo simulations and the decay time τ_{DA} as computed, for a non absorbing medium, as

$$\tau_{\text{DA}} = \frac{(L_0 + 2z_e)^2}{\pi D_{\text{DA}}}. \quad (6.2)$$

As opposed to the previous case, however, the τ/τ_{DA} ratio can evidently take values both above and below 1, depending subtly on the scattering anisotropy and the refractive index contrast of the slab. This observation might explain why retrieving the diffusion coefficient from a decay time measurement using equation (6.2) is sometimes regarded as a poor estimation, since this can lead to both over- or underestimated values [72]. This is further illustrated in the lower panel of Figure 6.1b for a couple of representative cases exhibiting opposite deviations that can persist even at higher optical thicknesses. This behavior is particularly interesting considering that, to date, experimental data and theoretical predictions are inconsistent. While the former bring generally evidences suggesting that retrieving D_{DA} through a decay time measurement would lead to a decreasing diffusion coefficient with decreasing thickness [87, 203], the latter have so far mainly provided arguments in favor of the opposite behavior [72, 85, 204, 205]. In this respect, our simulations show that there is a region in the parameter space where the τ/τ_{DA} ratio exceeds 1, which can lead to the experimentally observed decreasing diffusion coefficient with decreasing thickness. The analysis on the decay times confirms the importance of an accurate and precise modeling of the index contrast, which we think has been often overlooked, for example by considering a symmetric averaged contrast to model asymmetric experimental configurations [72, 85, 204].

6.2 Effective random-walk statistics

In order to explain the origin of the observed deviations for the in-plane transport, we focus on three significant configurations (highlighted as filled symbols in Figure 6.1a) representing key points of the observed peak for $g = 0$, i.e. $n = 1, 1.016$ and 1.1 . These three particular configurations were further investigated to collect detailed statistics at long times, with 10^{14} , 0.5×10^{14} and 10^{13} energy packets each. Performing simulations of such unprecedented magnitude required the use of the improved statistical sampling described in the previous Chapter and in Appendix B in order to accurately generate and represent the large number of random variates involved in the simulations.

As suggested by equation (6.1), the most straightforward insight on the diffusion coefficient D is obtained by directly looking at the distribution of the step-lengths performed during the random walk. In principle, each trajectory is generated according to the same exponential step-length distribution $p(\ell) = \exp(-\ell/l_s)/l_s$ (cfr. equation (2.89)). However, we find that the finite thickness of the slab configuration induces a positive correlation between a long permanence inside the sample and a higher




Fig. 6.2.: Late-time modification of the step length and scattering angle distributions for an optically thin slab with $OT = 1$, $g = 0$ and $n = 1.016$ and 1.1 . (a) Probability distribution of step lengths between consecutive scattering events performed by those energy packets that are transmitted at $t = 90 \text{ ps}$. The retrieved distributions exhibit heavier tails than the nominal one (dashed line). (b) Scattering angles become unevenly sampled at late times as well, exhibiting pronounced backwards and forward peaks.

probability of drawing longer step lengths. Figure 6.2 shows the histograms of the step lengths and scattering angles between two consecutive scattering events for those energy packets that were transmitted at $t = 90 \text{ ps}$ (corresponding to a path length of $\approx 27L_0$) compared with their nominal distributions implemented in the Monte Carlo algorithm (dashed lines). The SLD (Figure 6.2a) exhibits enhanced tails for all the three simulated refractive index contrasts, consistently with the observed enhancement of the diffusion rate (cfr. equation (6.1)). In this thin slab geometry, the nominal step length distribution provided by the pseudo-random number generator is sampled unevenly in such a way that all its moments are significantly modified: despite the fact that a long step in a very thin sample will generally cause the packet to exit the slab, those few packets that happen to remain inside will be able to reach long surviving times without undergoing many scattering events. In the case of refractive index contrasts close to 1, the distribution of the step lengths features a selective enhancement of the longer values, which is slightly more marked for $n = 1.016$. This is due to the fact that, even for such a small refractive index contrast, total internal reflection is already significant ($\theta_c = 79.8^\circ$). If internal reflections are absent, extremely narrow angular conditions must hold in order for the packet not to exit the slab. Conversely, even a tiny contrast allows to largely relax such condition, introducing a significant increase in the survival probability of a long-stepping energy packet while only marginally affecting others. In short, there is a positive correlation between long steps and shallow incidence angles, whose effects become apparent when such angles are the only ones undergoing total internal reflection (which also explains why the enhancement shown in Figure 6.1a is asymmetric around $n = 1$). On the other hand, with increasing contrast, more energy packets will be held inside the slab irrespective of their incidence angle (and hence of the length of their step), thus weakening the observed enhancement in the MSW growth rate.




Fig. 6.3.: (a) Time evolution of the ratio $\langle \ell^2 \rangle / 2 \langle \ell \rangle$ appearing in equation (6.1) and of $\langle \cos \theta \rangle$ resulting from the simulations. Each point is obtained considering only the energy packets transmitted within the corresponding time bin. Dashed lines represent the nominal values for the two distributions. (b) Time evolution of the SLD for $n = 1$ for energy packets transmitted at $t = 10, 20, 30, 40, 50, 60, 70, 80$ and 90 ps. Gray and purple curves show respectively the histogram of the step lengths drawn through the PRNG and of the steps taken inside the sample. The two only differ for the last step of each trajectory.

Interestingly, the sampling of the angular variables is also modified at late times, as shown in Figure 6.2b for the same set of simulations. While tracing each random trajectory, the cosines of the scattering (polar) angles θ are generated uniformly in $[-1, 1]$ through the pseudo-random number generator. On the contrary, the observed asymptotic $\cos \theta$ distribution exhibits two peaks for backwards and forward scattering. This can be intuitively understood by considering the fact that typical steps in a very long trajectory will be mostly aligned with the slab plane. As such, scattering angles close to $\theta = 0^\circ$ or 180° guarantee that the trajectory will continue within the slab irrespective of what azimuthal angle is drawn. Actually, since a typical step will not be in general perfectly parallel to the interfaces, a scattering angle of $\theta \approx 180^\circ$ should provide higher chances of staying inside the sample, hence its higher probability. This results in a $\cos \theta$ distribution with a slightly negative average value (Figure 6.3a, left axis), which also plays a role in determining the effective diffusion properties exhibited by the sample.

With reference to equation (6.1), we plot the quantity $\langle \ell^2 \rangle / 2 \langle \ell \rangle$ in Figure 6.3a (right axis), along with its nominal value of 1 (dashed line). At long times, each curve seems to saturate to an asymptotic value, suggesting the existence of a well-defined effective diffusion coefficient. The random-walk based picture of diffusion as expressed by equation (6.1) is qualitatively supported by the fact that also this figure of merit is enhanced for $n = 1.016$ (red curve), in accordance with Figure 6.1a. In principle, the overall diffusion process will be influenced by both the modified step-length and angular statistics, which in the investigated configurations appear to have opposite effects, as also shown in Figure 6.3a. While the latter would indeed tend to slightly slow down diffusion, the predominant effect is coming from the step lengths being substantially increased, leading to the observed enhanced in-plane diffusion especially for $n = 1.016$. Notably, different configurations might lead to a different balance between these two




Fig. 6.4.: (a) MSW expansion for the transmittance of a slab with ($n = 1.5$, $g = 0$, $OT^{-1} = 0.1$). As already discussed, in-plane propagation is slightly enhanced with respect to the DA $4D_{DA}t$ prediction. This deviation is apparently not supported by the asymptotic (b) step-length distribution (exhibiting a slightly lighter long-step tail than expected) and (c) scattering angle distribution (favoring backwards angles), that we approximate here as the effective distributions after a delay $t = 15\tau$.

effects, which also appear to saturate to their respective asymptotic values on slightly different time scales, further illustrating the need for additional investigations even for the simple model of a homogeneous and isotropic single slab.

The asymptotic nature of the effective diffusion coefficient in a thin slab is further highlighted in Figure 6.3b, where the time evolution of the SLD is shown for $n = 1$ (the $n = 1.016$ and 1.1 cases are analogous). The time-resolved distributions seem to converge towards a single asymptotic envelope distribution with a well-defined asymptotic decay rate which seems to be uniquely determined by the properties of the sample. It is interesting to compare the histogram of the actual steps performed inside the sample (blue curves) with the histogram of the ones drawn through the PRNG (gray curves). The two differ only for the last step, whose length is respectively considered either partially (up to the interface) or totally. At late times the two sets of curves become indistinguishable since, as expected, the contribution of the last step to the whole trajectory becomes eventually statistically negligible.

As a result of the transport statistics being directly altered by the sample configuration, an optically thin sample generally appears to be less scattering than it actually is. In other words, once the diffusive regime is reached, energy packets propagate as if scatterers were further apart than they actually are, i.e., with an *effective* transport mean free path greater than the one *intrinsic* to the material. Indeed, because of the asymptotic nature of these effects, only a small fraction of the incoming light is actually subject to this effective transport mean free path when studying thin samples. Yet, the effect is largely accessible experimentally [219] and, as we have seen in Section 4.2, similar deviations are can in fact be found even in more turbid media. As a matter of fact, the asymptotic nature of this effective transport regime makes it even more relevant from an experimental point of view, since it is often believed that the standard DA becomes progressively safer to apply as later times becomes accessible. Moreover, other applications can be envisioned




Fig. 6.5.: Time evolution of the ratio $\langle \ell^2 \rangle / 2 \langle \ell \rangle$ and of $\langle \cos \theta \rangle$ for a slab sample with ($n = 1.5$, $g = 0$, $OT^{-1} = 0.1$). Both figures converge asymptotically to a value that is lower than the nominal value of the simulation (dashed lines).

where multiple scattering in thin layers or confined geometries, even if limited to a very small fraction of incident light, could play a significant role [220, 221].

Albeit smaller, it is interesting to investigate whether the discrepancies that we discussed in Chapter 4 for the homogeneous slab arise from the same mechanism that we have described, which would suggest that modifications to the transport statistics can be still appreciable in more turbid media. At the same time, it is also interesting to determine how these deviations are gradually suppressed going towards optically thicker media. To this purpose, we performed a new simulation of a slab sample with $l_s = l'_s = 0.1$ mm and $L_0 = 1$ mm ($OT = 10$), using a refractive index contrast of $n = 1.5$ similar to that of our experimental slab sample (see Section 4.2). As we have seen in Chapter 5, the MSW expansion rate of such sample is slightly higher than that predicted by the DA. However, inspecting the step-length and scattering angle distribution reveals an apparently contradictory picture. As Figure 6.4 shows, while the MSW grows at a rate that is appreciably faster (of the order of %) of that expected, this does not reflect in the asymptotic ($t = 15\tau$) step-length nor in the scattering-angle distributions. As a matter of fact, the former is only slightly modified with respect to the nominal exponential, perhaps even exhibiting a *lower* probability of taking long steps, while the latter is slightly biased towards backwards scattering angles, which also goes in the direction of a smaller diffusion rate. This behavior is further confirmed by looking at the full evolution of the $\langle \ell^2 \rangle / 2 \langle \ell \rangle$ and $\langle \cos \theta \rangle$ descriptors (Figure 6.5), showing clearly that the asymptotic average step length and scattering angle are both smaller than expected. This interestingly shows how subtle deviations are still clearly present even at an optical thickness of $OT = 10$ where the DA is commonly employed, and that they are persistent at extremely long delays. In addition to this, the simulated data reveal a further layer of complexity in the effective transport regime that arises in bounded media, which we will address in the following.

6.3 A walk on the wild side of diffusion

In the following subsections, we briefly discuss few preliminary results regarding unexpected properties of transport in bounded media, which became apparent thanks to the magnitude of the simulations that we performed to calculate the exact solution to the transport problem.

6.3.1 Anisotropic transport in isotropic media

As we have seen in the previous Section, our investigation on the random-walk statistics arising in the slab geometry revealed a subtle interplay occurring between the actual thickness of the slab, the refractive index contrast and the scattering anisotropy, determining a transport regime that is basically diffusive on long time scales but which cannot be described in terms of the simple diffusive approximation. A different and asymptotic regime naturally emerges from the overall optical and geometric boundary conditions of the sample, and is univocally determined by them through yet unknown relations. In this respect, our findings recall a recently published work where it is analogously demonstrated that the link between microscopic (i.e., the scattering coefficient) and macroscopic (i.e., the diffusion coefficient) transport parameters remains unknown for diffusive anisotropic media [152]. Analogously, our results show that this link should be further investigated even in the isotropic case, especially for weakly scattering media. In particular, concerning microscopic optical properties such as g or l'_s , it seems appropriate to introduce a distinction between an intrinsic and an effective counterpart, where the former is the one that we are typically interested in retrieving while the latter might have a very different value and nature (e.g., tensorial instead of scalar) depending on incidental geometric conditions.

Driven by this observation, we performed a new analysis of our dataset, taking into account the translation symmetry of the problem along the xy -plane of the slab, collecting separate statistics on the step-length components taken along the different axes. In order to do so, it is convenient to derive the nominal probability density function (PDF) for the step components. This can be easily done in the simple isotropic ($g = 0$) case, where angles are sampled uniformly on a sphere. This allows to consider interchangeably the (fixed) reference frame of the slab and the (rotating) reference of the simulated energy packet. A generic step-length component will be therefore given by

$$\ell_i = \ell \cos \theta, \quad \text{with } i = x, y, z \quad (6.3)$$

which enables us to derive the probability density function $p(\ell_i)$ as that of the product of the two independent random variables ℓ and $\cos \theta$.

In general, the cumulative distribution function (CDF) for the product variable $X = UV$ can be written as

$$P(UV \leq x) = \int P(UV \leq x | U = u) p_U(u) du = \int P(uV \leq x) p_U(u) du$$




Fig. 6.6.: Comparison between derived and simulated probability density functions of (a) step lengths and (b) squared step lengths and their components as obtained from a single trajectory of 10^{10} steps in an infinitely extended volume.

$$= \int P\left(V \leq \frac{x}{u}\right) p_U(u) du = \int P_V\left(\frac{x}{u}\right) p_U(u) du. \quad (6.4)$$

A convenient choice is that of using the cumulative distribution of the exponential distribution (2.60), which is non-null only for $\ell > 0$, and the PDF of the uniform distribution, which has a limited support. By plugging their expressions into equation (6.4), following integration by parts eventually we obtain

$$\begin{aligned} P(X \leq x) &= \int_0^1 1 - \exp(-\mu_s x/u) du = 1 - \int_0^1 \exp(-\mu_s x/u) du \\ &= 1 - \exp(-\mu_s x) + \mu_s x \Gamma(0, \mu_s x) \end{aligned} \quad (6.5)$$

where $\Gamma(0, x)$ is the incomplete gamma function (also reported as the exponential integral function $E_1(x)$). The probability density function is then easily found as the derivative

$$p(x) = \frac{dP_X(x)}{dx} = \mu_s \Gamma(0, \mu_s x). \quad (6.6)$$

Once we have obtained the PDF, the probability functions for the squared variables is also simply found by exploiting the fact that the step lengths and the length of its components are positive quantities. By doing so, we can generally write

$$p_{\xi^2}(x) = \frac{d}{dx} P(\xi^2 \leq x) = \frac{d}{dx} P(\xi \leq \sqrt{x}) = \frac{dP_\xi}{d\sqrt{x}} = p(\sqrt{x}) \frac{1}{2\sqrt{x}}, \quad (6.7)$$

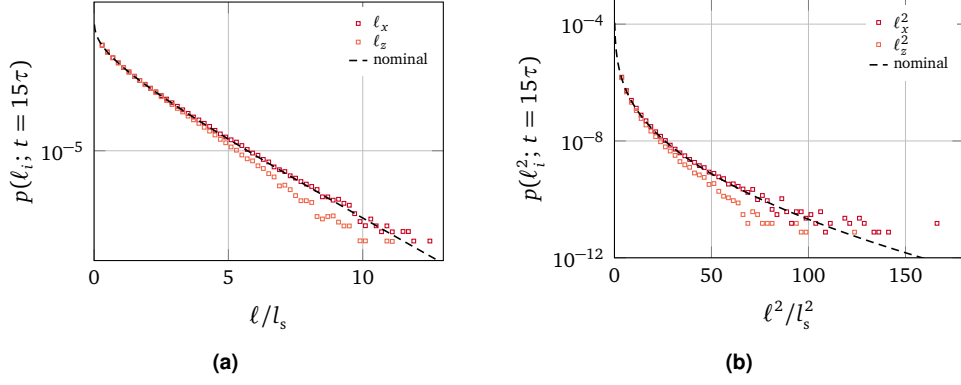


Fig. 6.7.: (a) Decomposition of the SLD shown in Figure 6.4b into its x and z components, revealing how they are differently affected by the presence of the plane boundary. The expression (6.6) for an unbounded medium is shown for comparison. Due to internal reflections, at long times also the z -component SLD continues smoothly beyond L_0 . (b) Decomposition of the squared-step-length distribution, compared with equation (6.9), showing a qualitatively similar behavior.

which gives

$$p_{\ell^2}(x) = \frac{\mu_s}{2\sqrt{x}} \exp(-\mu_s \sqrt{x}) \quad (6.8)$$

$$p_{\ell_i^2}(x) = \frac{\mu_s}{2\sqrt{x}} \Gamma(0, -\mu_s \sqrt{x}) \quad (6.9)$$

respectively for ℓ^2 and ℓ_i^2 . The relevant distributions that we obtained along with their first moments are listed in the following table:

X	$p(x)$	$\langle x \rangle$
ℓ	$\exp(-\ell/l_s)/l_s$	l_s
ℓ^2	$\frac{1}{2l_s\sqrt{\ell}} \exp(-\ell/l_s)$	$2l_s^2$
ℓ_i	$\Gamma(0, \ell/l_s)/l_s$	$l_s/2$
ℓ_i^2	$\frac{1}{2l_s\sqrt{\ell}} \Gamma(0, \sqrt{\ell}/l_s)$	$2l_s^2/3$

A comparison between these expressions and the step-length distributions of an isotropic trajectory in an unbounded medium is shown in Figure 6.6.

By analyzing the late-time step-length distribution of our $OT = 10$ simulation in terms of their x and z components (y is statistically equivalent to x), an interesting picture emerges. Figure 6.7 illustrates the difference between the in-plane and axial components of the steps of those packets transmitted at long times. Despite the fact that the simulated system is isotropic and isotropically scattering, an *anisotropic* transport regime establishes at long times characterized by a constrained step-length distribution along z . However, also the step lengths along the plane are modified, as can be more clearly appreciated from Figure 6.8 where we have plotted the $\langle \ell_i^2 \rangle / 2 \langle \ell_i \rangle$ ratio. This unexpected over-compensation for the decreased ‘diffusivity’ along z is what determines




Fig. 6.8.: Time evolution of the (a) anisotropic $\langle \ell_i^2 \rangle / 2 \langle \ell_i \rangle$ ratios, revealing how an enhanced in-plane MSW expansion is compatible with an overall negative anisotropy factor g and smaller probability of taking long steps, as shown in Figure 6.5. In-plane and axial components undergo opposite modifications, with a more marked suppression along z accompanied by a net increase along x and y . (b) These modifications do not seem to affect the shape of the transmitted profile, whose excess kurtosis γ_2 still vanishes with time.

the enhanced MSW growth that we have found in Figures 6.4 and 5.3. Nonetheless, the overall effect is still that of a decreased probability of taking long steps, due to the more marked reduction of $\langle \ell_z^2 \rangle / 2 \langle \ell_z \rangle$. Notably, the onset of this anisotropic transport regime does not seem to affect the shape of the transmitted profiles, whose excess kurtosis converges normally to 0 (Figure 6.8b).

As expected, even more pronounced deviations can be found in the previously presented $OT = 1$ configurations, yet with a qualitatively similar behavior (Figure 6.9), where the z component is particularly suppressed due to the low refractive index contrast. In such a confined configuration, in fact, long trajectories will be composed of steps mainly lying in the plane of the slab, with the extreme case of matched refractive index contrast where no step can be taken having an axial component larger than the thickness of the slab.

At this stage, to our knowledge, there is no analytical model available to describe how this anisotropic transport regime sets in depending on the optical thickness and the refractive index contrast of the slab. In particular, even though this regime is largely determined by the presence and type of boundaries, its effect is fundamentally different from apparently similar boundary effects described in the literature [75, 76], which can be usually taken into account through some refined extrapolated boundary conditions. This cannot be the case here, since extrapolated boundary conditions correct significantly quantities such as the total transmittance which, conversely, would be negligibly affected by asymptotic modifications of the effective diffusion coefficient, especially in thicker media. This does not mean that the effect that we described cannot be accessed experimentally. On the contrary, we can now interpret the small discrepancy measured for our homogeneous sample in Chapter 4 as a direct experimental evidence of the onset of this anisotropic regime. State-of-the-art detectors used in time-resolved configuration are able to access 8 decades of dynamic range [219], corresponding to

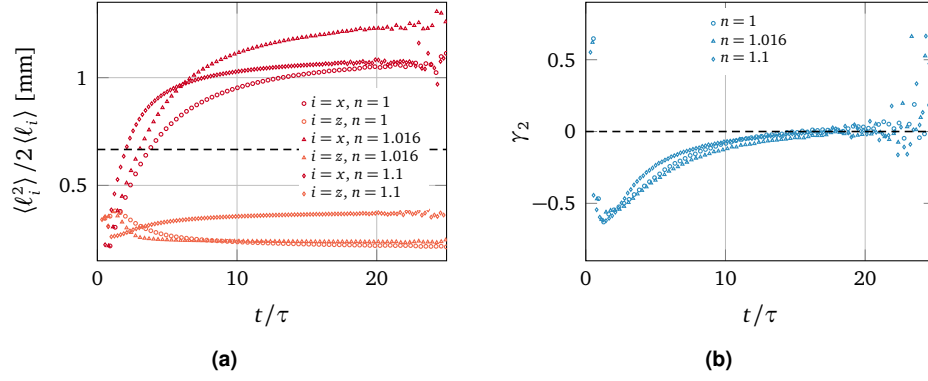


Fig. 6.9.: Time evolution of the (a) anisotropic $\langle \ell_i^2 \rangle / 2 \langle \ell_i \rangle$ ratio for the different refractive index contrasts of the OT = 1 configurations. (b) Despite the much larger deviations observed, also in this case the excess kurtosis seems to converge towards 0 without any finite offset.

a suppression greater than $\sim e^{18}$, a range within which even our weakly scattering configurations have almost reached their asymptotic regime.

6.3.2 Anomalous diffusion in homogeneous media

In the previous Chapters, we have often used the mean-square width (i.e., the second moment of the time-resolved profiles) as a convenient descriptor of transport. However, as we have seen in Section 4.3 for a few extreme cases, only the profile shape conveys the full information for a proper characterization. An interesting trade-off which has been recently considered in the literature of continuous-time random walks (CTRWs), consists of analyzing the spatio-temporal evolution of the profiles at different (continuous) moments [222, 223]. As a matter of fact, many transport processes exhibit moments of displacements with asymptotic behavior

$$\langle |x|^q \rangle \sim t^{q\nu(q)}, \quad (6.10)$$

and the study of $q\nu(q)$ as a function of q provides a more complete characterization of the process than does the single scalar $2\nu(2)$. In standard diffusive and anomalous diffusive processes, the function $\nu(q)$ is actually a constant (i.e., $\nu = 1/2$ in the normal diffusive case, $\nu > 1/2$ for superdiffusion and $\nu = 1$ for ballistic propagation). However, characterizing the dichotomy between normal and anomalous diffusion based only on the value of $\nu(2)$ fails at providing the full picture [223], and one should look further than $\langle |x|^2 \rangle(t)$. Ideally, the complete information is represented by the Green's function or propagator of the process. Often this is not possible to obtain exactly, though asymptotic methods can be exploited to provide useful approximations in the form of similarity solutions. These are expressed, in the limit of $t \rightarrow \infty$ by the self-similar collapse relation

$$p(x, t) \approx t^{-1/\nu} \mathcal{P}\left(\frac{x}{t^{1/\nu}}\right). \quad (6.11)$$

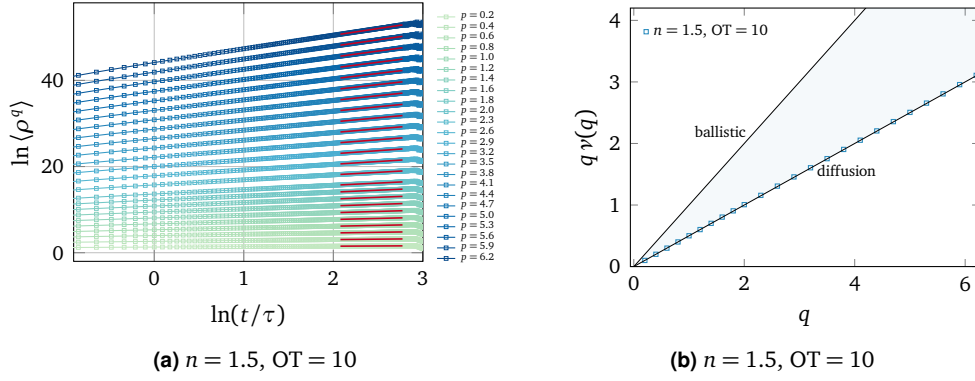


Fig. 6.10.: (a) Time evolution of moments $q \in [0.2, 6.2]$ calculated for the transmitted profiles of the $n = 1.5$, $OT = 10$ sample. A linear fit of $\ln \langle \rho^q \rangle$ versus $\ln \frac{t}{\tau}$ is performed in the range between $8 < t/\tau < 16$ where the asymptotic power-law growth is established. (b) The obtained exponents reveal a strongly self-similar and diffusive regime, as expected for a turbid, homogeneous medium. Ballistic ($q\nu(q) = 1$) and diffusive ($q\nu(q) = 1/2$) lines are drawn as guides to the eye.

In the case of normal diffusion $\nu = 2$ and \mathcal{P} is a Gaussian, while in the superdiffusive case there are examples in which \mathcal{P} is a Lévy density [224]. The approximation (6.11) is valid only in the central scaling region (CSR), i.e., excluding the non-scaling tails of the distribution which are populated by ‘particles’ that have undergone exceptionally large displacements. The concept of *strong* self-similarity can be introduced to identify those systems where $\nu(q) = 1/\nu$, $\forall q$, as it is the case for normal diffusion and Lévy flights. On the contrary, *weak* self-similarity refers to propagation dynamics giving rise to non-trivial $\nu(q)$ functions. Typically, the small- q part of $\nu(q)$ (which passes through the origin of the $(\nu(q), q)$ plane) refers to the central scaling region of the asymptotic profile containing the majority of the ‘particles’, while the large- q range is determined by relatively few tail individuals which have traveled exceptionally far. For this reason, as confirmed by analytical, numerical and experimental evidence, in most known systems $\nu(q)q$ is a piece-wise linear function experiencing a change of slope at a certain value of $q = q_c$ [222, 223, 225, 226]. This is the typical behavior exhibited, for instance, by Lévy walks [227, 228]. For most systems studied in the literature, the large- q branch is therefore characterized by a ballistic growth ($\nu(q) = 1$) which sets in for those high moments that are eventually dominated by the few particles that propagate ballistically. On the contrary, the presence of ballistically expanding tails does not necessarily imply any form of weak self-similarity. The simple telegrapher model is one such example that is both normally diffusive and strongly self-similar: its tail structure is not self-similar at all times, but its contribution to the profile is mild enough that as $t \rightarrow \infty$ all moments are still determined by the CSR. As a final note, the value of q_c at which the slope of weakly self-similar systems can change depending on the parameters of the transport process involved. If this occurs before $q = 2$, measuring the expansion rate of the MSW would actually probe the non-scaling tails rather than the density in the CSR. In other words, a system might still exhibit a perfectly linear $2\nu(2) = 1$ while being both weakly self-similar and anomalously diffusive. Therefore, a full characterization beyond




Fig. 6.11.: Time evolution of moments $q \in [0.2, 6.2]$ calculated for the transmitted profiles of the optically thin samples ($OT = 1$) with refractive index contrast values of (a) $n = 1$, (b) $n = 1.016$ and (c) $n = 1.1$. A linear fit of $\ln \langle \rho^q \rangle$ versus $\ln \frac{t}{\tau}$ is performed in the range between $8 < t/\tau < 16$ where the asymptotic power-law growth is established.

the simple second moment is always recommended to correctly identify the nature of transport.

In the following, we present a preliminary discussion of the numerical results that we obtained based on the previous simulations on slab geometry systems, considering the asymptotic behavior for different moments of the transmitted profiles using $|x| = \rho$. A first, relevant test case is that of the $OT = 10$ slab that we used as an analogue of our experimental sample of Section 4.2. As we have seen, this system is characterized by an asymptotic transport regime that is slightly anisotropic, but still diffusive. This is confirmed by the late- t analysis of the moments as shown in Figure 6.10. In the analysis, the power-law exponents for each moment can be more accurately obtained by performing a linear fit of $\ln \langle \rho^q \rangle$ versus $\ln \frac{t}{\tau}$, considering only late times between $\ln 8 \leq \ln \frac{t}{\tau} \leq \ln 16$ where statistical noise is negligible and transport has already reached its asymptotic regime. As expected, analysis of the fractional moments confirms both the diffusive and the strongly self-similar nature of light transport in the turbid slab, with all exponents lying on the analytic diffusive line $\langle \rho^q \rangle = t^{q/2}$.

The situation is more interesting when we turn to our optically thin configurations, which we have simulated with further increased statistics. In this case, plotting the expansion of the profiles for different widths reveals that the asymptotic regime is reached at even later times, as can also be appreciated by the slow convergence of the step-length distributions to their asymptotic form (see, for instance, Figure 6.2). To this purpose, the power-law exponents have been retrieved in a range of delays between $\ln 16 \leq \ln \frac{t}{\tau} \leq \ln 20$, beyond which simulation noise starts to be appreciable (Figure 6.11). In particular, when plotting the full behavior of $q \nu(q)$ we can actually appreciate that none of the sample is truly diffusive, and that slight deviations in the linearity that are barely appreciable by considering just the second moment (Figure 6.1a) actually fit into a weakly superdiffusive pattern, as can be appreciated in Figure 6.12a. This result is quite surprising given that the velocities of packets in the simulation is constant and the variance of the SLD is finite and is probably due to the time-dependent nature of the SLD.

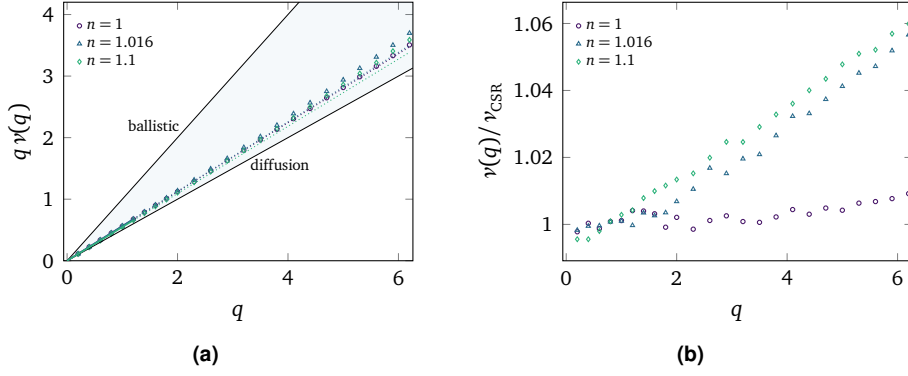


Fig. 6.12.: (a) Fractional exponents obtained by the linear fits of Figure 6.11, revealing a superdiffusive characteristic of transport in the optically-thin, low-contrast slab configuration. Error bars (not shown), calculated as the 95 % confidence intervals returned by a linear least square method, are approximately the same size of the symbols. A linear fit in the small- q range $[0.2, 1.2]$ reveals a linear superdiffusive behavior with $\nu(q \leq 1.2) = \nu_{\text{CSR}} = 0.5605(3)$, $0.5644(3)$ and $0.5461(7)$ respectively for $n = 1$, 1.016 and 1.1 . Fitted lines are extrapolated to higher moments as guides to the eye. (b) Relative residuals between the values of $q\nu(q)$ obtained from the simulations and the extrapolated linear scaling $\nu_{\text{CSR}}(n)$, highlighting the presence of weakly self-similar anomalous diffusion in the presence of a mismatch refractive index.

We have fitted the fractional exponents $q \in [0.2, 1.2]$ of Figure 6.12a to characterize the scaling properties of the CSR, obtaining a good linear agreement. Notably, even in the small- q range a small superdiffusive behavior is to be found for all refractive index contrasts, with $\nu(q \leq 1.2) = \nu_{\text{CSR}} = 0.5605(3)$, $0.5644(3)$ and $0.5461(7)$ respectively for $n = 1$, 1.016 and 1.1 .

By analyzing the moments in terms of their self-similarity, however, differences between the configurations emerge. In Figure 6.12b, we plot the relative deviation between the line $q\nu_{\text{CSR}}(n)$ extrapolated to higher values of q and the simulated data, highlighting a clear qualitative difference arising in the presence of a refractive-index contrast. In particular, while the index-matched configuration seems to exhibit a strongly self-similar (and superdiffusive) behavior, the other configurations show linearly increasing relative deviations. This is not only a hallmark of weak self-similarity, but is also signaling that the commonly studied piece-wise linear type of weak self-similarity does not apply to this case. On the contrary, a second order dependence on q seems also to be taken into account.

As we discussed in the introduction, the piece-wise model of weakly self-similar diffusion refers just to the simple case where the ballistic tails of the propagator eventually represent the prevailing contribution in the determination of higher moments. However, this regime is unlikely to be relevant for our samples. The transmitted profiles recorded at $16 \leq t/\tau \leq 20$ certainly do not exhibit any ballistic peak on the tails, and basically all energy packets transmitted within that time range have experienced at least a few scattering events, given the large but finite magnitude of our simulations. This might explain why the piece-wise linear model fails at describing this configuration, which

apparently requires the need of a superlinear term. As a matter of fact, it is believed that weak self-similarity might occur in several forms, of which the piece-wise linear model is only one simple case. Nonetheless, the literature on general, non piece-wise linear examples of weakly self-similar diffusion is extremely limited, and focused on very different models with no losses, power-law SLDs and propagation velocities drawn from a distribution with a time-growing variance [229]. It is possible that the time-varying characteristic (which, in our case, regards the SLD rather than the velocities) is the key element giving rise to more general weakly self-similar dynamics. In this respect, light transport in a simple homogeneous, isotropic scattering slab might represent an interesting physical platform for the experimental study of a broader array of transport regimes than expected. To this end, an optimal trade-off should be sought after between the refractive index contrast, the absolute and the relative optical thickness of the sample, in order to design an experiment where a weakly self-similar transport regime is reached compatibly with the instrumental sensitivity and quantitative accuracy available for spatio-temporal imaging techniques. In fact, as of now, the slow convergence to the asymptotic transport regime exhibited in optically thin systems represents a challenging issue, even from the point of view of numerical simulations.

Now that the study of transport phenomena has grown to such a mature research field spanning over many different branches of science, it is perhaps surprising that such a rich transport physics, comprising important and complex concepts as anisotropy, weakly self-similar scaling and anomalous diffusion, can all be found and studied in such a simple and explored model as the single plane-parallel homogeneous slab with isotropic, annealed disorder.

Bibliography

1. Einstein, A. Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen. *Annalen der Physik* **322**, 549–560 (1905).
2. Gonzalez, M. C., Hidalgo, C. A. & Barabasi, A.-L. Understanding individual human mobility patterns. *Nature* **453**, 779–782 (2008).
3. Van den Akker, J. Scattering and absorption of light in paper and other diffusing media. *Tappi* **32**, 498–501 (1949).
4. Thomas, G. E. & Stamnes, K. *Radiative transfer in the atmosphere and ocean* (Cambridge University Press, 2002).
5. Marshak, A. & Davis, A. *3D radiative transfer in cloudy atmospheres* (Springer Science & Business Media, 2005).
6. Tuchin, V. *Tissue optics: light scattering methods and instruments for medical diagnosis* (SPIE press Bellingham, 2007).
7. Martelli, F., Del Bianco, S., Ismaelli, A. & Zaccanti, G. *Light Propagation through Biological Tissue and Other Diffusive Media: Theory, Solutions and Software* (SPIE Press, Bellingham, 2010).
8. Sato, H., Fehler, M. C. & Maeda, T. *Seismic wave propagation and scattering in the heterogeneous Earth* (Springer, 2012).
9. Campillo, M. & Paul, A. Long-range correlations in the diffuse seismic coda. *Science* **299**, 547–549 (2003).
10. Hansen, J. E. & Hovenier, J. Interpretation of the polarization of Venus. *Journal of the Atmospheric Sciences* **31**, 1137–1160 (1974).
11. Porco, C. C. *et al.* Simulations of the dynamical and light-scattering behavior of Saturn's rings and the derivation of ring particle and disk properties. *The Astronomical Journal* **136**, 2172 (2008).
12. Barthelemy, P., Bertolotti, J. & Wiersma, D. S. A Lévy flight for light. *Nature* **453**, 495–498 (2008).
13. Segev, M., Silberberg, Y. & Christodoulides, D. N. Anderson localization of light. *Nature Photonics* **7**, 197–204 (2013).
14. Caruso, F., Crespi, A., Ciriolo, A. G., Sciarrino, F. & Osellame, R. Fast escape of a quantum walker from an integrated photonic maze. *Nature Communications* **7** (2016).
15. Froufe-Pérez, L. S. *et al.* Role of Short-Range Order and Hyperuniformity in the Formation of Band Gaps in Disordered Photonic Materials. *Physical Review Letters* **117**, 053902 (2016).

16. Pauly, M., Kollig, T. & Keller, A. *Metropolis light transport for participating media* (Springer, 2000).
17. Schmidt, T.-W., Pellacini, F., Nowrouzezahrai, D., Jarosz, W. & Dachsbacher, C. *State of the art in artistic editing of appearance, lighting and material in Computer Graphics Forum* (2015).
18. Planck, M. *The theory of heat radiation* (P. Blakiston's Son & Co., Philadelphia, 1914).
19. Chandrasekhar, S. *Radiative transfer* (Oxford, Clarendon Press, 1950).
20. Schuster, A. Radiation through a foggy atmosphere. *The Astrophysical Journal* **21**, 1 (1905).
21. Mishchenko, M. I. Vector radiative transfer equation for arbitrarily shaped and arbitrarily oriented particles: a microphysical derivation from statistical electromagnetics. *Applied Optics* **41**, 7114–7134 (2002).
22. Ripoll, J. Derivation of the scalar radiative transfer equation from energy conservation of Maxwell's equations in the far field. *JOSA A* **28**, 1765–1775 (2011).
23. Ishimaru, A. *Wave propagation and scattering in random media* (Academic press New York, 1978).
24. Born, M. & Wolf, E. *Principles of optics: electromagnetic theory of propagation, interference and diffraction of light* (CUP Archive, 2000).
25. Van de Hulst, H. C. *Multiple light scattering: tables, formulas, and applications* (Elsevier, 2012).
27. Henyey, L. G. & Greenstein, J. L. Diffuse radiation in the galaxy. *The Astrophysical Journal* **93**, 70–83 (1941).
28. Graaff, R., Aarnoudse, J. G., de Mul, F. F. & Jentink, H. W. Similarity relations for anisotropic scattering in absorbing media. *Optical Engineering* **32**, 244–252 (1993).
29. Bertolotti, J. *et al.* Engineering disorder in superdiffusive Levy glasses. *Advanced Functional Materials* **20**, 965–968 (2010).
30. Mackowski, D. W. & Mishchenko, M. I. Calculation of the T matrix and the scattering matrix for ensembles of spheres. *JOSA A* **13**, 2266–2278 (1996).
31. Feinberg, J. Self-pumped, continuous-wave phase conjugator using internal reflection. *Optics Letters* **7**, 486–488 (1982).
32. Yaqoob, Z., Psaltis, D., Feld, M. S. & Yang, C. Optical phase conjugation for turbidity suppression in biological samples. *Nature Photonics* **2**, 110–115 (2008).
33. Vellekoop, I., Lagendijk, A. & Mosk, A. Exploiting disorder for perfect focusing. *Nature Photonics* **4**, 320–322 (2010).
34. Popoff, S., Goetschy, A., Liew, S., Stone, A. D. & Cao, H. Coherent control of total transmission of light through disordered media. *Physical Review Letters* **112**, 133903 (2014).
35. Redding, B., Liew, S. F., Sarma, R. & Cao, H. Compact spectrometer based on a disordered photonic chip. *Nature Photonics* **7**, 746–751 (2013).

36. Goorden, S. A., Horstmann, M., Mosk, A. P., Škorić, B. & Pinkse, P. W. Quantum-secure authentication of a physical unclonable key. *Optica* **1**, 421–424 (2014).
37. Van Albada, M. P. & Lagendijk, A. Observation of weak localization of light in a random medium. *Physical Review Letters* **55**, 2692 (1985).
38. Wiersma, D. S., Bartolini, P., Lagendijk, A. & Righini, R. Localization of light in a disordered medium. *Nature* **390**, 671–673 (1997).
39. Prum, R. O., Dufresne, E. R., Quinn, T. & Waters, K. Development of colour-producing β -keratin nanostructures in avian feather barbs. *Journal of The Royal Society Interface* **6**, S253–S265 (2009).
40. Muller, N., Haberkorn, J., Marichy, C. & Scheffold, F. Silicon hyperuniform disordered photonic materials with a pronounced gap in the shortwave infrared. *Advanced Optical Materials* **2**, 115–119 (2014).
41. Leseur, O., Pierrat, R. & Carminati, R. High-density hyperuniform materials can be transparent. *Optica* **3**, 763–767 (2016).
42. Badon, A. *et al.* Spatio-temporal imaging of light transport in highly scattering media under white light illumination. *Optica* **3**, 1160–1166 (2016).
43. Wyman, D. R., Patterson, M. S. & Wilson, B. C. Similarity relations for the interaction parameters in radiation transport. *Applied Optics* **28**, 5243–5249 (1989).
44. Contini, D., Martelli, F. & Zaccanti, G. Photon migration through a turbid slab described by a model based on diffusion approximation. I. Theory. *Applied Optics* **36**, 4587–4599 (1997).
45. Vynck, K., Pierrat, R. & Carminati, R. Polarization and spatial coherence of electromagnetic waves in uncorrelated disordered media. *Physical Review A* **89**, 013842 (2014).
46. Roux, L., Mareschal, P., Vukadinovic, N., Thibaud, J.-B. & Greffet, J.-J. Scattering by a slab containing randomly located cylinders: comparison between radiative transfer and electromagnetic simulation. *JOSA A* **18**, 374–384 (2001).
47. Fraden, S. & Maret, G. Multiple light scattering from concentrated, interacting suspensions. *Physical Review Letters* **65**, 512 (1990).
48. Bressel, L. & Reich, O. Theoretical and experimental study of the diffuse transmission of light through highly concentrated absorbing and scattering materials: Part I: Monte-Carlo simulations. *Journal of Quantitative Spectroscopy and Radiative Transfer* **146**, 190–198 (2014).
49. Modest, M. F. *Radiative heat transfer* (Academic Press, 2013).
50. Machida, M., Panasyuk, G. Y., Schotland, J. C. & Markel, V. A. The Green's function for the radiative transport equation in the slab geometry. *Journal of Physics A: Mathematical and Theoretical* **43**, 065402 (2010).
51. Liemert, A. & Kienle, A. Exact and efficient solution of the radiative transport equation for the semi-infinite medium. *Scientific Reports* **3** (2013).
52. Wilson, B. & Adam, G. A Monte Carlo model for the absorption and flux distributions of light in tissue. *Medical Physics* **10**, 824–830 (1983).

53. Wang, L., Jacques, S. L. & Zheng, L. MCML—Monte Carlo modeling of light transport in multi-layered tissues. *Computer Methods and Programs in Biomedicine* **47**, 131–146 (1995).
54. Alerstam, E., Svensson, T. & Andersson-Engels, S. Parallel computing with graphics processing units for high-speed Monte Carlo simulation of photon migration. *Journal of Biomedical Optics* **13**, 060504–060504 (2008).
55. Fang, Q. Mesh-based Monte Carlo method using fast ray-tracing in Plücker coordinates. *Biomedical Optics Express* **1**, 165 (2010).
56. Patterson, M. S., Chance, B. & Wilson, B. C. Time resolved reflectance and transmittance for the noninvasive measurement of tissue optical properties. *Applied Optics* **28**, 2331–2336 (1989).
57. Bassani, M., Martelli, F., Zaccanti, G. & Contini, D. Independence of the diffusion coefficient from absorption: experimental and numerical evidence. *Optics Letters* **22**, 853–855 (1997).
58. Durduran, T., Yodh, A., Chance, B. & Boas, D. Does the photon-diffusion coefficient depend on absorption? *JOSA A* **14**, 3358–3365 (1997).
59. Pierrat, R., Greffet, J.-J. & Carminati, R. Photon diffusion coefficient in scattering and absorbing media. *JOSA A* **23**, 1106–1110 (2006).
60. Alerstam, E. *Optical spectroscopy of turbid media: time-domain measurements and accelerated monte carlo modelling*. PhD thesis (Lund University, 2011).
61. Haskell, R. C. *et al.* Boundary conditions for the diffusion equation in radiative transfer. *JOSA A* **11**, 2727–2741 (1994).
62. Diebold, M. P. *Application of Light Scattering to Coatings: A User's Guide* (Springer, 2014).
63. Comsa, D., Farrell, T. & Patterson, M. Quantification of bioluminescence images of point source objects using diffusion theory models. *Physics in Medicine and Biology* **51**, 3733 (2006).
64. Martelli, F. *et al.* Hybrid heuristic time dependent solution of the radiative transfer equation for the slab in *European Conference on Biomedical Optics* (2009), 7369_1B.
65. Morse, P. M. & Feshbach, H. *Methods of theoretical physics* (McGraw-Hill New York, 1953).
66. Martelli, F., Sassaroli, A., Zaccanti, G. & Yamada, Y. Properties of the light emerging from a diffusive medium: angular dependence and flux at the external boundary. *Physics in Medicine and Biology* **44**, 1257 (1999).
67. MacKintosh, F. & John, S. Diffusing-wave spectroscopy and multiple scattering of light in correlated random media. *Physical Review B* **40**, 2383 (1989).
68. Yoo, K., Liu, F. & Alfano, R. When does the diffusion approximation fail to describe photon transport in random media? *Physical Review Letters* **64**, 2647 (1990).
69. Yoo, K. & Alfano, R. Time-resolved coherent and incoherent components of forward light scattering in random media. *Optics Letters* **15**, 320–322 (1990).

70. Lemieux, P.-A., Vera, M. & Durian, D. Diffusing-light spectroscopies beyond the diffusion limit: The role of ballistic transport and anisotropic scattering. *Physical Review E* **57**, 4498 (1998).
71. Elaloufi, R., Carminati, R. & Greffet, J.-J. Time-dependent transport through scattering media: from radiative transfer to diffusion. *Journal of Optics A: Pure and Applied Optics* **4**, S103 (2002).
72. Elaloufi, R., Carminati, R. & Greffet, J.-J. Diffusive-to-ballistic transition in dynamic light transmission through thin scattering slabs: a radiative transfer approach. *JOSA A* **21**, 1430–1437 (2004).
73. Kienle, A. & Patterson, M. S. Improved solutions of the steady-state and the time-resolved diffusion equations for reflectance from a semi-infinite turbid medium. *JOSA A* **14**, 246–254 (1997).
74. Venugopalan, V., You, J. & Tromberg, B. Radiative transport in the diffusion approximation: an extension for highly absorbing media and small source-detector separations. *Physical Review E* **58**, 2395 (1998).
75. Chen, N. G. & Bai, J. Monte Carlo approach to modeling of boundary conditions for the diffusion equation. *Physical Review Letters* **80**, 5321 (1998).
76. Popescu, G., Mujat, C. & Dogariu, A. Evidence of scattering anisotropy effects on boundary conditions of the diffusion equation. *Physical Review E* **61**, 4523 (2000).
77. Garofalakis, A. *et al.* Characterization of the reduced scattering coefficient for optically thin samples: theory and experiments. *Journal of Optics A: Pure and Applied Optics* **6**, 725 (2004).
78. Farrell, T. J., Patterson, M. S. & Wilson, B. A diffusion theory model of spatially resolved, steady-state diffuse reflectance for the noninvasive determination of tissue optical properties in vivo. *Medical Physics* **19**, 879–888 (1992).
79. Martelli, F., Bassani, M., Alianelli, L., Zangheri, L. & Zaccanti, G. Accuracy of the diffusion equation to describe photon migration through an infinite medium: numerical and experimental investigation. *Physics in Medicine and Biology* **45**, 1359 (2000).
80. Jacques, S. L. Optical properties of biological tissues: a review. *Physics in Medicine and Biology* **58**, R37 (2013).
81. Durian, D. Influence of boundary reflection and refraction on diffusive photon transport. *Physical Review E* **50**, 857 (1994).
82. Furutsu, K. & Yamada, Y. Diffusion approximation for a dissipative random medium and the applications. *Physical Review E* **50**, 3634 (1994).
83. Ripoll, J., Yessayan, D., Zacharakis, G. & Ntziachristos, V. Experimental determination of photon propagation in highly absorbing and scattering media. *JOSA A* **22**, 546–551 (2005).
84. Zhang, Z. *et al.* Wave transport in random media: The ballistic to diffusive transition. *Physical Review E* **60**, 4843 (1999).
85. Zhang, X. & Zhang, Z.-Q. Wave transport through thin slabs of random media with internal reflection: Ballistic to diffusive transition. *Physical Review E* **66**, 016612 (2002).

86. Yaroshevsky, A., Glasser, Z., Granot, E. & Sternklar, S. Transition from the ballistic to the diffusive regime in a turbid medium. *Optics Letters* **36**, 1395–1397 (2011).
87. Kop, R. H., de Vries, P., Sprik, R. & Lagendijk, A. Observation of anomalous transport of strongly multiple scattered light in thin disordered slabs. *Physical Review Letters* **79**, 4369 (1997).
88. Swartling, J. *Biomedical and atmospheric applications of optical spectroscopy in scattering media*. PhD thesis (Department of Physics, Lund University, 2002).
89. Grabtchak, S. & Whelan, W. M. Separation of absorption and scattering properties of turbid media using relative spectrally resolved cw radiance measurements. *Biomedical Optics Express* **3**, 2371–2380 (2012).
90. Groenhuis, R., Ferwerda, H. A. & Ten Bosch, J. Scattering and absorption of turbid materials determined from reflection measurements. 1: Theory. *Applied Optics* **22**, 2456–2462 (1983).
91. Hammer, M., Roggan, A., Schweitzer, D. & Muller, G. Optical properties of ocular fundus tissues-an in vitro study using the double-integrating-sphere technique and inverse Monte Carlo simulation. *Physics in Medicine and Biology* **40**, 963 (1995).
92. Pickering, J. W. *et al.* Double-integrating-sphere system for measuring the optical properties of tissue. *Applied Optics* **32**, 399–410 (1993).
93. Dam, J. S., Dalgaard, T., Fabricius, P. E. & Andersson-Engels, S. Multiple polynomial regression method for determination of biomedical optical properties from integrating sphere measurements. *Applied Optics* **39**, 1202–1209 (2000).
94. Dimofte, A., Finlay, J. C. & Zhu, T. C. A method for determination of the absorption and scattering properties interstitially in turbid media. *Physics in Medicine and Biology* **50**, 2291 (2005).
95. Hall, G., Jacques, S. L., Eliceiri, K. W. & Campagnola, P. J. Goniometric measurements of thick tissue using Monte Carlo simulations to obtain the single scattering anisotropy coefficient. *Biomedical Optics Express* **3**, 2707–2719 (2012).
96. Drezek, R., Dunn, A. & Richards-Kortum, R. Light scattering from cells: finite-difference time-domain simulations and goniometric measurements. *Applied Optics* **38**, 3651–3661 (1999).
97. Kienle, A. *et al.* Spatially resolved absolute diffuse reflectance measurements for noninvasive determination of the optical scattering and absorption coefficients of biological tissue. *Applied Optics* **35**, 2304–2314 (1996).
98. Cuccia, D. J., Bevilacqua, F., Durkin, A. J., Ayers, F. R. & Tromberg, B. J. Quantitation and mapping of tissue optical properties using modulated imaging. *Journal of Biomedical Optics* **14**, 024012–024012 (2009).
99. Scheffold, F. & Block, I. D. Rapid high resolution imaging of diffusive properties in turbid media. *Optics Express* **20**, 192–200 (2012).
100. Fishkin, J. B. & Gratton, E. Propagation of photon-density waves in strongly scattering media containing an absorbing semi-infinite plane bounded by a straight edge. *JOSA A* **10**, 127–140 (1993).

101. Tromberg, B. J., Svaasand, L. O., Tsay, T.-T. & Haskell, R. C. Properties of photon density waves in multiple-scattering media. *Applied Optics* **32**, 607–616 (1993).
102. Svensson, T. *et al.* Exploiting breakdown of the similarity relation for diffuse light transport: simultaneous retrieval of scattering anisotropy and diffusion constant. *Optics Letters* **38**, 437–439 (2013).
103. Stott, J. & Boas, D. A. A practical comparison between time-domain and frequency-domain diffusive optical imaging systems in *Biomedical Topical Meetings* **71** (2002), 626–628.
104. Dam, J. S., Yavari, N., Sørensen, S. & Andersson-Engels, S. Real-time absorption and scattering characterization of slab-shaped turbid samples obtained by a combination of angular and spatially resolved measurements. *Applied Optics* **44**, 4281–4290 (2005).
105. Alerstam, E. & Svensson, T. Observation of anisotropic diffusion of light in compacted granular porous materials. *Physical Review E* **85**, 040301 (2012).
106. Torricelli, A. *et al.* Time domain functional NIRS imaging for human brain mapping. *Neuroimage* **85**, 28–50 (2014).
107. Svensson, T., Andersson-Engels, S., Einarsdóttir, M. & Svanberg, K. In vivo optical characterization of human prostate tissue using near-infrared time-resolved spectroscopy. *Journal of Biomedical Optics* **12**, 014022–014022 (2007).
108. Svensson, T., Alerstam, E., Johansson, J. & Andersson-Engels, S. Optical porosimetry and investigations of the porosity experienced by light interacting with porous media. *Optics Letters* **35**, 1740–1742 (2010).
109. D’Andrea, C. *et al.* Time-resolved optical imaging through turbid media using a fast data acquisition system based on a gated CCD camera. *Journal of Physics D: Applied Physics* **36**, 1675 (2003).
110. Sperling, T., Buehrer, W., Aegerter, C. M. & Maret, G. Direct determination of the transition to localization of light in three dimensions. *Nature Photonics* **7**, 48–52 (2012).
111. Delpy, D. T. *et al.* Estimation of optical pathlength through tissue from direct time of flight measurement. *Physics in Medicine and Biology* **33**, 1433 (1988).
112. Hebden, J. C., Kruger, R. A. & Wong, K. Time resolved imaging through a highly scattering medium. *Applied Optics* **30**, 788–794 (1991).
113. Hebden, J. C., Arridge, S. R. & Delpy, D. T. Optical imaging in medicine: I. Experimental techniques. *Physics in Medicine and Biology* **42**, 825 (1997).
114. Dunsby, C. & French, P. Techniques for depth-resolved imaging through turbid media including coherence-gated imaging. *Journal of Physics D: Applied Physics* **36**, R207 (2003).
115. Shimizu, K., Ishimaru, A., Reynolds, L. & Bruckner, A. P. Backscattering of a picosecond pulse from densely distributed scatterers. *Applied Optics* **18**, 3484–3488 (1979).
116. Wang, L., Ho, P., Liu, C., Zhang, G. & Alfano, R. Ballistic 2-D imaging through scattering walls using an ultrafast optical Kerr gate. *Science* **253**, 769–771 (1991).
117. Duncan, M. D., Mahon, R., Tankersley, L. L. & Reintjes, J. Time-gated imaging through scattering media using stimulated Raman amplification. *Optics Letters* **16**, 1868–1870 (1991).

118. Shah, J. Ultrafast luminescence spectroscopy using sum frequency generation. *IEEE Journal of Quantum Electronics* **24**, 276–288 (1988).
119. Yodh, A., Kaplan, P. & Pine, D. Pulsed diffusing-wave spectroscopy: High resolution through nonlinear optical gating. *Physical Review B* **42**, 4744 (1990).
120. Faris, G. W. & Banks, M. Upconverting time gate for imaging through highly scattering media. *Optics Letters* **19**, 1813–1815 (1994).
121. Le Tolguenec, G., Lantz, E. & Devaux, F. Imaging through scattering media by parametric amplification of images: study of the resolution and the signal-to-noise ratio. *Applied Optics* **36**, 8292–8297 (1997).
122. Yoo, K., Xing, Q. & Alfano, R. Imaging objects hidden in highly scattering media using femtosecond second-harmonic-generation cross-correlation time gating. *Optics Letters* **16**, 1019–1021 (1991).
123. Wiersma, D. S., Muzzi, A., Colocci, M. & Righini, R. Time-resolved experiments on light diffusion in anisotropic random media. *Physical Review E* **62**, 6681 (2000).
124. Trebino, R. Second-harmonic-generation rings and refractive-index measurement in uniaxial crystals. *Applied Optics* **20**, 2090–2096 (1981).
125. Zhang, D., Kong, Y. & Zhang, J.-y. Optical parametric properties of 532-nm-pumped beta-barium-borate near the infrared absorption edge. *Optics Communications* **184**, 485–491 (2000).
126. Wasylczyk, P. Ultracompact autocorrelator for femtosecond laser pulses. *Review of Scientific Instruments* **72**, 2221–2223 (2001).
127. Haus, H. A. Mode-locking of lasers. *IEEE Journal of Selected Topics in Quantum Electronics* **6**, 1173–1185 (2000).
129. Savo, R. *Transport in complex heterogeneous photonic structures*. PhD thesis (Università di Firenze, 2013).
130. Trebino, R. *Frequency-Resolved Optical Gating: The Measurement of Ultrashort Laser Pulses* (Springer, 2000).
131. Burrese, M. *et al.* Bright-White Beetle Scales Optimise Multiple Scattering of Light. *Scientific Reports* **4** (2014).
132. Potenza, M. A. *et al.* Three dimensional imaging of short pulses. *Optics Communications* **229**, 381–390 (2004).
133. Firester, A. Image upconversion: part III. *Journal of Applied Physics* **41**, 703–709 (1970).
134. Chiou, W. Geometric optics theory of parametric image upconversion. *Journal of Applied Physics* **42**, 1985–1993 (1971).
135. Devaux, F. & Lantz, E. Ultrahigh-speed imaging by parametric image amplification. *Optics Communications* **118**, 25–27 (1995).
136. Abraham, E. *et al.* Real-time two-dimensional imaging in scattering media by use of a femtosecond Cr^{4+} : forsterite laser. *Optics Letters* **25**, 929–931 (2000).

137. Lantz, E. & Devaux, F. Parametric amplification of images: from time gating to noiseless amplification. *Selected Topics in Quantum Electronics, IEEE Journal of* **14**, 635–647 (2008).
138. Pedersen, C., Karamehmedović, E., Dam, J. S. & Tidemand-Lichtenberg, P. Enhanced 2D-image upconversion using solid-state lasers. *Optics Express* **17**, 20885–20890 (2009).
139. Bassi, A. *et al.* Time-gated optical projection tomography. *Optics Letters* **35**, 2732–2734 (2010).
140. Dam, J. S., Pedersen, C. & Tidemand-Lichtenberg, P. High-resolution two-dimensional image upconversion of incoherent light. *Optics Letters* **35**, 3796–3798 (2010).
142. Vos, W. L., Tukker, T. W., Mosk, A. P., Lagendijk, A. & IJzerman, W. L. Broadband mean free path of diffuse light in polydisperse ensembles of scatterers for white light-emitting diode lighting. *Applied Optics* **52**, 2602–2609 (2013).
144. Alerstam, E., Svensson, T. & Andersson-Engels, S. *CUDAMCML-User manual and implementation notes* (2009).
146. Milsom, P. A ray-optic, Monte Carlo, description of a Gaussian beam waist–applied to reverse saturable absorption. *Applied Physics B* **70**, 593–599 (2000).
147. Alerstam, E., Andersson-Engels, S. & Svensson, T. Improved accuracy in time-resolved diffuse reflectance spectroscopy. *Optics Express* **16**, 10440–10454 (2008).
148. Hu, H., Strybulevych, A., Page, J. H., Skipetrov, S. E. & van Tiggelen, B. A. Localization of ultrasound in a three-dimensional elastic network. *Nature Physics* **4**, 945–948 (2008).
149. Cherroret, N., Skipetrov, S. & Van Tiggelen, B. Transverse confinement of waves in three-dimensional random media. *Physical Review E* **82**, 056603 (2010).
150. Bouchard, J.-P. *et al.* Reference optical phantoms for diffuse optical spectroscopy. Part 1–Error analysis of a time resolved transmittance characterization method. *Optics Express* **18**, 11495–11507 (2010).
151. Comelli, D. *et al.* In vivo time-resolved reflectance spectroscopy of the human forehead. *Applied Optics* **46**, 1717–1725 (2007).
152. Alerstam, E. Anisotropic diffusive transport: Connecting microscopic scattering and macroscopic transport properties. *Physical Review E* **89**, 063202 (2014).
153. Marchesini, R., Bertoni, A., Andreola, S., Melloni, E. & Sichirollo, A. Extinction and absorption coefficients and scattering phase functions of human tissues in vitro. *Applied Optics* **28**, 2318–2324 (1989).
154. Premože, S., Ashikhmin, M., Tessendorf, J., Ramamoorthi, R. & Nayar, S. *Practical rendering of multiple scattering effects in participating media* in *Proceedings of the Fifteenth Eurographics conference on Rendering Techniques* (2004), 363–374.
155. Nilsson, A. M., Stureson, C., Liu, D. L. & Andersson-Engels, S. Changes in spectral shape of tissue optical properties in conjunction with laser-induced thermotherapy. *Applied Optics* **37**, 1256–1267 (1998).
156. Mourant, J. R. *et al.* Mechanisms of light scattering from biological cells relevant to noninvasive optical-tissue diagnostics. *Applied Optics* **37**, 3586–3593 (1998).

157. Pattelli, L., Savo, R., Burrelli, M. & Wiersma, D. Spatio-temporal visualization of light transport in complex photonic structures. *Light: Science & Applications* **5**, e16090 (2016).
158. Mazzamuto, G., Pattelli, L., Toninelli, C. & Wiersma, D. S. Deducing effective light transport parameters in optically thin systems. *New Journal of Physics* **18**, 023036 (2016).
159. Kienle, A. Anisotropic light diffusion: an oxymoron? *Physical Review Letters* **98**, 218104 (2007).
160. Cortese, L. *et al.* Anisotropic Light Transport in White Beetle Scales. *Advanced Optical Materials* **3**, 1337–1341 (2015).
161. Svensson, T. *et al.* Light diffusion in quenched disorder: Role of step correlations. *Physical Review E* **89**, 022141 (2014).
162. Bruck, R. *et al.* Device-level characterization of the flow of light in integrated photonic circuits using ultrafast photomodulation spectroscopy. *Nature Photonics* **9**, 54–60 (2015).
163. Liemert, A. & Kienle, A. Analytical solution of the radiative transfer equation for infinite-space fluence. *Physical Review A* **83**, 015804 (2011).
164. Liemert, A. & Kienle, A. Infinite space Green's function of the time-dependent radiative transfer equation. *Biomedical Optics Express* **3**, 543–551 (2012).
165. Graaff, R. *et al.* Condensed Monte Carlo simulations for the description of light transport. *Applied Optics* **32**, 426–434 (1993).
166. Kienle, A. & Patterson, M. S. Determination of the optical properties of turbid media from a single Monte Carlo simulation. *Physics in Medicine and Biology* **41**, 2221 (1996).
167. Pifferi, A., Taroni, P., Valentini, G. & Andersson-Engels, S. Real-time method for fitting time-resolved reflectance and transmittance measurements with a Monte Carlo model. *Applied Optics* **37**, 2774–2780 (1998).
168. Thueler, P. *et al.* In vivo endoscopic tissue diagnostics based on spectroscopic absorption, scattering, and phase function properties. *Journal of Biomedical Optics* **8**, 495–503 (2003).
169. Palmer, G. M. & Ramanujam, N. Monte Carlo-based inverse model for calculating tissue optical properties. Part I: Theory and validation on synthetic phantoms. *Applied Optics* **45**, 1062–1071 (2006).
170. Xu, H., Farrell, T. J. & Patterson, M. S. Investigation of light propagation models to determine the optical properties of tissue from interstitial frequency domain fluence measurements. *Journal of Biomedical Optics* **11**, 041104–041104 (2006).
171. Alerstam, E., Andersson-Engels, S. & Svensson, T. White Monte Carlo for time-resolved photon migration. *Journal of Biomedical Optics* **13**, 041304–041304 (2008).
172. Martinelli, M. *et al.* Analysis of single Monte Carlo methods for prediction of reflectance from turbid media. *Optics Express* **19**, 19627–19642 (2011).
173. Leonardi, L. & Burns, D. H. Quantitative measurements in scattering media: Photon time-of-flight analysis with analytical descriptors. *Applied Spectroscopy* **53**, 628–636 (1999).

174. Pfefer, T. J. *et al.* Reflectance-based determination of optical properties in highly attenuating tissue. *Journal of Biomedical Optics* **8**, 206–215 (2003).
175. Chauchard, F. *et al.* Least-squares support vector machines modelization for time-resolved spectroscopy. *Applied Optics* **44**, 7091–7097 (2005).
176. Alerstam, E. *et al.* Single-fiber diffuse optical time-of-flight spectroscopy. *Optics Letters* **37**, 2877–2879 (2012).
177. Jäger, M., Foschum, F. & Kienle, A. Application of multiple artificial neural networks for the determination of the optical properties of turbid media. *Journal of Biomedical Optics* **18**, 057005–057005 (2013).
178. Rajaram, N., Nguyen, T. H. & Tunnell, J. W. Lookup table-based inverse model for determining optical properties of turbid media. *Journal of Biomedical Optics* **13**, 050501–050501 (2008).
179. Nilsson, A. M., Berg, R. & Andersson-Engels, S. Measurements of the optical properties of tissue in conjunction with photodynamic therapy. *Applied Optics* **34**, 4609–4619 (1995).
180. Roggan, A., Albrecht, H. J., Doerschel, K., Minet, O. & Mueller, G. J. *Experimental set-up and Monte-Carlo model for the determination of optical tissue properties in the wavelength range 330 to 1100 nm in International Symposium on Biomedical Optics Europe'94* (1995), 21–36.
181. Karlsson, H., Fredriksson, I., Larsson, M. & Strömberg, T. Inverse Monte Carlo for estimation of scattering and absorption in liquid optical phantoms. *Optics Express* **20**, 12233–12246 (2012).
182. Hennessy, R., Lim, S. L., Markey, M. K. & Tunnell, J. W. Monte Carlo lookup table-based inverse model for extracting optical properties from tissue-simulating phantoms using diffuse reflectance spectroscopy. *Journal of Biomedical Optics* **18**, 037003–037003 (2013).
183. Simpson, C. R., Kohl, M., Essenpreis, M. & Cope, M. Near-infrared optical properties of ex vivo human skin and subcutaneous tissues measured using the Monte Carlo inversion technique. *Physics in Medicine and Biology* **43**, 2465 (1998).
185. Alerstam, E. *et al.* Next-generation acceleration and code optimization for light transport in turbid media using GPUs. *Biomedical Optics Express* **1**, 658–675 (2010).
186. Ramella-Roman, J., Prahl, S. & Jacques, S. Three Monte Carlo programs of polarized light transport into scattering media: part I. *Optics Express* **13**, 4420–4438 (2005).
187. Fang, Q. & Boas, D. A. Monte Carlo simulation of photon migration in 3D turbid media accelerated by graphics processing units. *Optics Express* **17**, 20178 (2009).
188. Doronin, A. & Meglinski, I. Online object oriented Monte Carlo computational tool for the needs of biomedical optics. *Biomedical Optics Express* **2**, 2461 (2011).
189. Doering, C. R., Ray, T. S. & Glasser, M. L. Long transmission times for transport through a weakly scattering slab. *Physical Review A* **45**, 825 (1992).
190. Boguñá, M., Porra, J. M. & Masoliver, J. Persistent random walk model for transport through thin slabs. *Physical Review E* **59**, 6517 (1999).

191. Duparré, A. & Kassam, S. Relation between light scattering and the microstructure of optical thin films. *Applied Optics* **32**, 5475–5480 (1993).
192. Perelman, L. *et al.* Observation of periodic fine structure in reflectance from biological tissue: a new technique for measuring nuclear size distribution. *Physical Review Letters* **80**, 627 (1998).
193. De Bruin, D. M. *et al.* Optical phantoms of varying geometry based on thin building blocks with controlled optical properties. *Journal of Biomedical Optics* **15**, 025001–025001 (2010).
194. Leonetti, M. & López, C. Measurement of transport mean-free path of light in thin systems. *Optics Letters* **36**, 2824–2826 (2011).
195. Chan, E. K. *et al.* Effects of compression on soft tissue optical properties. *Selected Topics in Quantum Electronics, IEEE Journal of* **2**, 943–950 (1996).
196. Backman, V. *et al.* Polarized light scattering spectroscopy for quantitative measurement of epithelial cellular structures in situ. *Selected Topics in Quantum Electronics, IEEE Journal of* **5**, 1019–1026 (1999).
197. Du, Y. *et al.* Optical properties of porcine skin dermis between 900 nm and 1500 nm. *Physics in Medicine and Biology* **46**, 167 (2001).
198. Ugryumova, N., Matcher, S. J. & Attenburrow, D. P. Measurement of bone mineral density via light scattering. *Physics in Medicine and Biology* **49**, 469 (2004).
199. Darling, C. L., Huynh, G. D. & Fried, D. Light scattering properties of natural and artificially demineralized dental enamel at 1310nm. *Journal of Biomedical Optics* **11**, 034023–034023 (2006).
200. Pattelli, L., Mazzamuto, G., Wiersma, D. S. & Toninelli, C. Diffusive light transport in semitransparent media. *Physical Review A* **94**, 043846 (2016).
201. Douglass, K. M. & Dogariu, A. Measuring diffusion coefficients independently of boundary conditions. *Optics Letters* **34**, 3379–3381 (2009).
202. Bolin, F. P., Preuss, L. E., Taylor, R. C. & Ference, R. J. Refractive index of some mammalian tissues using a fiber optic cladding method. *Applied Optics* **28**, 2297–2303 (1989).
203. Rivas, J. G., Sprik, R., Lagendijk, A., Noordam, L. & Rella, C. Static and dynamic transport of light close to the Anderson localization transition. *Physical Review E* **63**, 046613 (2001).
204. Ramakrishna, S. A. & Kumar, N. Diffusion of particles moving with constant speed. *Physical Review E* **60**, 1381–1389 (1999).
205. Gopal, V., Anantha Ramakrishna, S., Sood, A. & Kumar, N. Photon transport in thin disordered slabs. *Pramana-Journal of Physics* **56**, 767–778 (2001).
206. Spinelli, L. *et al.* Calibration of scattering and absorption properties of a liquid diffusive medium at NIR wavelengths. Time-resolved method. *Optics Express* **15**, 6589–6604 (2007).

207. Chance, B. *et al.* Comparison of time-resolved and -unresolved measurements of deoxyhemoglobin in brain. *Proceedings of the National Academy of Sciences* **85**, 4971–4975 (1988).
208. Wilson, B. C. & Jacques, S. L. Optical reflectance and transmittance of tissues: principles and applications. *IEEE Journal of Quantum Electronics* **26**, 2186–2199 (1990).
209. Svensson, T. *Pharmaceutical and biomedical applications of spectroscopy in the photon migration regime*. PhD thesis (Lund University, 2008).
210. Fishkin, J. B., Fantini, S., Gratton, E., *et al.* Gigahertz photon density waves in a turbid medium: theory and experiments. *Physical Review E* **53**, 2307 (1996).
211. Ntziachristos, V. & Chance, B. Accuracy limits in the determination of absolute optical properties using time-resolved NIR spectroscopy. *Medical Physics* **28**, 1115–1124 (2001).
212. Scheffold, F., Lenke, R., Tweer, R. & Maret, G. Localization or classical diffusion of light? *Nature* **398**, 206–207 (1999).
213. Wiersma, D. S., Rivas, J. G., Bartolini, P., Lagendijk, A. & Righini, R. Reply: Localization or classical diffusion of light? *Nature* **398**, 207 (1999).
215. Gómez-Medina, R. *et al.* Negative scattering asymmetry parameter for dipolar particles: Unusual reduction of the transport mean free path and radiation pressure. *Physical Review A* **85**, 035802 (2012).
216. Wald, A. A Reprint of "A Method of Estimating Plane Vulnerability Based on Damage of Survivors" by Abraham Wald (Operations Evaluation Group, Center for Naval Analyses, 1980).
217. Mangel, M. & Samaniego, F. J. Abraham Wald's work on aircraft survivability. *Journal of the American Statistical Association* **79**, 259–267 (1984).
218. Svensson, T. *et al.* Holey random walks: Optics of heterogeneous turbid composites. *Physical Review E* **87**, 022120 (2013).
219. Tosi, A. *et al.* Fast-gated single-photon counting technique widens dynamic range and speeds up acquisition time in time-resolved measurements. *Optics Express* **19**, 10735–10746 (2011).
220. Cao, H., Xu, J., Seelig, E. & Chang, R. P. Microlaser made of disordered media. *Applied Physics Letters* **76**, 2997–2999 (2000).
221. Wu, X. *et al.* Random lasing in weakly scattering systems. *Physical Review A* **74**, 053812 (2006).
222. Castiglione, P., Mazzino, A., Muratore-Ginanneschi, P. & Vulpiani, A. On strong anomalous diffusion. *Physica D: Nonlinear Phenomena* **134**, 75–93 (1999).
223. Ferrari, R., Manfroi, A. & Young, W. Strongly and weakly self-similar diffusion. *Physica D: Nonlinear Phenomena* **154**, 111–137 (2001).
224. Zumofen, G. & Klafter, J. Scale-invariant motion in intermittent chaotic systems. *Physical Review E* **47**, 851 (1993).
225. Gal, N. & Weihs, D. Experimental evidence of strong anomalous diffusion in living cells. *Physical Review E* **81**, 020903 (2010).

226. Metzler, R., Jeon, J.-H., Cherstvy, A. G. & Barkai, E. Anomalous diffusion models and their properties: non-stationarity, non-ergodicity, and ageing at the centenary of single particle tracking. *Physical Chemistry Chemical Physics* **16**, 24128–24164 (2014).
227. Allegrini, P. *et al.* Scaling breakdown: a signature of aging. *Physical Review E* **66**, 015101 (2002).
228. Bernabó, P., Burioni, R., Lepri, S. & Vezzani, A. Anomalous transmission and drifts in one-dimensional Lévy structures. *Chaos, Solitons & Fractals* **67**, 11–19 (2014).
229. Andersen, K., Castiglione, P., Mazzino, A. & Vulpiani, A. Simple stochastic models showing strong anomalous diffusion. *The European Physical Journal B-Condensed Matter and Complex Systems* **18**, 447–452 (2000).

Webpages

- @26. Prahl, S. *Mie Scattering calculator* <http://omlc.org/calc/mie_calc.html>.
- @128. Kärtner, F. *Ultrafast Optics* Massachusetts Institute of Technology: MIT OpenCourseWare. <<http://ocw.mit.edu>>. License: Creative Commons BY-NC-SA.
- @141. Thorlabs Inc. *UV-Curing Optical Adhesives* <http://www.thorlabs.de/newgrouppage9.cfm?objectgroup_id=196>.
- @143. Mazzamuto, G. & Pattelli, L. *MCPlusPlus Documentation* <<http://www.lens.unifi.it/quantum-nanophotonics/mcplusplus/>>.
- @145. Wikipedia. *Logistic distribution* — *Wikipedia, The Free Encyclopedia* <https://en.wikipedia.org/wiki/Logistic_distribution>.
- @184. Alerstam, E. *GPU Monte Carlo* <<http://www.atomic.physics.lu.se/biophotonics/research/monte-carlo-simulations/gpu-monte-carlo/>>.
- @214. Mazzamuto, G. & Pattelli, L. *LUT routine for the single slab geometry* <<http://www.lens.unifi.it/quantum-nanophotonics/mcplusplus/lut>>.

A.1 Derivation of the radiative transfer equation

In this appendix we review a possible derivation of the radiative transfer equation (RTE) using Poynting's theorem for energy conservation. Let us start considering the time-averaged expression (2.12) that we derived in section 2.1

$$\frac{1}{v} \frac{\partial \langle \mathbf{S}(\mathbf{r}) \rangle \cdot \mathbf{s}_j}{\partial t} + \left\langle \frac{dP_{\text{abs}}}{dV} \right\rangle (\mathbf{s} \cdot \mathbf{s}_j) + \mathbf{s}_j \cdot \nabla (\langle \mathbf{S}(\mathbf{r}) \rangle \cdot \mathbf{s}_j) = 0, \quad (\text{A.1})$$

which ensures that energy conservation is rotationally invariant and holds in any particular direction \mathbf{s}_j . This equation is still valid also in the presence of a time-varying dependence (i.e., a modulation of the source) on a time-scale that is slower than the electromagnetic oscillation ω . On the other hand, we must remember that it is based on the assumption that the electric and magnetic fields are mutually orthogonal, and therefore holds only in the far-field of the particles.

In order to obtain the RTE we will apply energy conservation to a small volume δV (cfr. Fig. 2.3) containing N particles. Integrating equation (A.1) over δV we obtain

$$\int_{\delta V} (\mathbf{s} \cdot \mathbf{s}_j) \left[\frac{1}{v} \frac{\partial}{\partial t} S(\mathbf{r} - \mathbf{r}') + \left\langle \frac{dP_{\text{abs}}}{dV}(\mathbf{r} - \mathbf{r}') \right\rangle + \mathbf{s}_j \cdot \nabla_{\mathbf{r}'} S(\mathbf{r} - \mathbf{r}') \right] d^3 r' = 0 \quad (\text{A.2})$$

where $S(\mathbf{r})$ is the magnitude of the time-averaged Poynting vector at \mathbf{r} .

Inside the sample volume δV we can decompose the energy flux $\langle \mathbf{S} \rangle$ as the sum of the contribution of a scattered flux $\langle \mathbf{S}_{\text{sc}} \rangle$ and an incident flux $\langle \mathbf{S}_{\text{inc}} \rangle$, where the latter is now a general function accounting for the incoming flux from outside δV . If the volume δV is negligible with respect to the volume $V - \delta V \gg \delta V$ that is responsible for $\langle \mathbf{S}_{\text{inc}} \rangle$, we can assume that

$$\int_{\delta V} \langle \mathbf{S}_{\text{inc}} \rangle \cdot \mathbf{s}_j dV \gg \int_{\delta V} \langle \mathbf{S}_{\text{sc}} \rangle \cdot \mathbf{s}_j dV \quad (\text{A.3})$$

which is an analogue to the *dipolar approximation* for multiple scattering of small particles, in that it neglects any self-induction term and assumes that the incident field on each particle is simply given by the sum of all scattered fields excluded its own. Analogously, given the small size of δV , we can make the assumption that the incident field at any point inside such sample volume can be approximated to the volume-averaged incident flow

$$\langle \mathbf{S}_{\text{inc}}(\mathbf{r}) \rangle \cdot \mathbf{s} \simeq \|\mathbf{S}_{\text{inc}}(\mathbf{r})\|_v \quad (\text{A.4})$$

where \mathbf{s} is the direction of the flow of energy at \mathbf{r} , as usual.

Using these approximations, the first and second term of equation (A.2) can be rewritten as

$$\frac{1}{v} \frac{\partial}{\partial t} \int_{\delta V} \langle \mathbf{S} \rangle \cdot \mathbf{s}_j dV \simeq \frac{1}{v} \frac{\partial}{\partial t} \|\mathbf{S}_{\text{inc}}\|_v w_r(\mathbf{s}_j) \delta V \quad (\text{A.5})$$

$$\int_{\delta V} \left\langle \frac{dP_{\text{abs}}}{dV} \right\rangle (\mathbf{s} \cdot \mathbf{s}_j) dV \simeq N \sigma_a \|\mathbf{S}_{\text{inc}}\|_v w_r(\mathbf{s}_j), \quad (\text{A.6})$$

in terms of the volume-averaged incident flux, where $w_r(\mathbf{s}_j)$ is the probability distribution of the energy flux to each angle defined in equation (2.39), N is the number of particles in δV and σ_a is their absorption cross-section (2.23).

On the other hand, the third term of equation (A.2) becomes

$$\begin{aligned} \int_{\delta V} (\mathbf{s} \cdot \mathbf{s}_j) \mathbf{s}_j \cdot \nabla_{r'} S(\mathbf{r} - \mathbf{r}') d^3 r' &= w_r(\mathbf{s}_j) \|\nabla_{r'} \mathbf{S}_{\text{inc}} \cdot \mathbf{s}_j\|_v \delta V \\ &+ w_r(\mathbf{s}_j) \int_{\delta V} \mathbf{s}_j \cdot \nabla_{r'} S_{\text{sc}}(\mathbf{r} - \mathbf{r}') d^3 r' \end{aligned} \quad (\text{A.7})$$

where we have introduced the volume-averaged change in $\langle \mathbf{S} \rangle$ as

$$\|\nabla_{r'} \mathbf{S}_{\text{inc}} \cdot \mathbf{s}_j\|_v = \frac{1}{\delta V} \int_{\delta V} \mathbf{s}_j \cdot \nabla_{r'} S_{\text{inc}}(\mathbf{r} - \mathbf{r}') d^3 r'. \quad (\text{A.8})$$

At this point, assuming that the far-field approximation is valid within δV (which we already used in order for the cross-sections to be additive in equation (A.6)) and yet that δV is small enough to consider $\mathbf{r} \simeq \mathbf{r}'$, we can interchange $\nabla_r \leftrightarrow \nabla_{r'}$ and write

$$\|\nabla \mathbf{S}_{\text{inc}} \cdot \mathbf{s}_j\|_v \simeq \mathbf{s}_j \cdot \nabla \|\mathbf{S}_{\text{inc}}\|_v. \quad (\text{A.9})$$

Plugging this substitution into equation (A.7) we obtain

$$\begin{aligned} \int_{\delta V} (\mathbf{s} \cdot \mathbf{s}_j) \mathbf{s}_j \cdot \nabla_{r'} S(\mathbf{r} - \mathbf{r}') d^3 r' &= \\ w_r(\mathbf{s}_j) \mathbf{s}_j \cdot \nabla \|\mathbf{S}_{\text{inc}}(\mathbf{r})\|_v \delta V &+ w_r(\mathbf{s}_j) \int_{\delta V} \mathbf{s}_j \cdot \nabla_{r'} S_{\text{sc}}(\mathbf{r} - \mathbf{r}') d^3 r' \end{aligned} \quad (\text{A.10})$$

where the last term is that accounting for the scattered flux, both outgoing and inwards from outside δV . We can apply Gauss' theorem provided that both these terms are properly represented

$$\int_{\delta V} \mathbf{s}_j \cdot \nabla_{r'} S_{\text{sc}}(\mathbf{r} - \mathbf{r}') d^3 r' = \int_{\Sigma} S_{\text{sc}}^{(\text{out})}(\mathbf{r} - \mathbf{r}') \mathbf{s}_j \cdot \mathbf{s}' dS' - \int_{\Sigma} S_{\text{sc}}^{(\text{in})}(\mathbf{r} - \mathbf{r}') \mathbf{s}_j \cdot \mathbf{s}' dS' \quad (\text{A.11})$$

where Σ is the surface enclosing δV . The first term can be written as the sum of the contributions from the N scattering particles in δV , while the inward flux represents the contribution from the whole outer volume $V - \delta V$. To obtain an expression for

such inward flux, we can exploit the fact that we have a medium with average optical properties which are homogeneous at every position. We can therefore assume that at each boundary between neighboring volume elements δV , the total inward and outward flux can be considered equivalent. For each point \mathbf{r} lying on the surface Σ we can then write that

$$\int_{\Sigma} S_{\text{sc}}^{(\text{in})}(\mathbf{r} - \mathbf{r}') \mathbf{s}_j \cdot \mathbf{s}' dS' \simeq \sum_{i=1}^N \int_{\Sigma} (\langle \mathbf{S}_{\text{sc}}(\mathbf{r}) \rangle_i \cdot \mathbf{s}_i) \mathbf{s}_j \cdot \mathbf{s}' dS' \quad (\text{A.12})$$

where $\langle \mathbf{S}_{\text{sc}}(\mathbf{r}) \rangle_i$ is the energy flux scattered by the i -th particle inside δV , given by equation (2.36). Making use of the approximation (A.4) we obtain

$$\begin{aligned} \int_{\Sigma} S_{\text{sc}}^{(\text{in})}(\mathbf{r} - \mathbf{r}') \mathbf{s}_j \cdot \mathbf{s}' dS' &\simeq \sigma_{\text{tot}} \sum_{i=1}^N \|\mathbf{S}_{\text{inc}}(\mathbf{r})\|_{\text{v}} \int_{\Sigma} w_{\mathbf{r}}(\mathbf{s}_i) \frac{p(\mathbf{s}_i, \mathbf{s}_j)}{|\mathbf{r}' - \mathbf{r}_i|^2} dS' \\ &\simeq N \sigma_{\text{tot}} \|\mathbf{S}_{\text{inc}}(\mathbf{r})\|_{\text{v}} \int_{4\pi} w_{\mathbf{r}}(\mathbf{s}') p(\mathbf{s}', \mathbf{s}_j) d\Omega', \end{aligned} \quad (\text{A.13})$$

where in the last step we have identified $dS'/|\mathbf{r} - \mathbf{r}_i|^2$ as the solid angle $d\Omega$, since the total flux traversing Σ will not depend on the position of the particles within δV .

At this point we have obtained the expressions for all the terms appearing in equation (A.2), which can be rewritten as

$$\begin{aligned} \frac{1}{v} \frac{\partial}{\partial t} \|\mathbf{S}_{\text{inc}}\|_{\text{v}} w_{\mathbf{r}}(\mathbf{s}_j) \delta V + N \sigma_{\text{a}} \|\mathbf{S}_{\text{inc}}\|_{\text{v}} w_{\mathbf{r}}(\mathbf{s}_j) + \mathbf{s}_j \cdot \nabla \|\mathbf{S}_{\text{inc}}\|_{\text{v}} w_{\mathbf{r}}(\mathbf{s}_j) \delta V \\ + N \sigma_{\text{s}} \|\mathbf{S}_{\text{inc}}\|_{\text{v}} w_{\mathbf{r}}(\mathbf{s}_j) - N \sigma_{\text{tot}} \int_{4\pi} \|\mathbf{S}_{\text{inc}}\|_{\text{v}} w_{\mathbf{r}}(\mathbf{s}') p(\mathbf{s}', \mathbf{s}_j) d\Omega' \end{aligned} \quad (\text{A.14})$$

and normalized by δV to remove the dependence from the arbitrary averaging volume

$$\begin{aligned} \frac{1}{v} \frac{\partial}{\partial t} \|\mathbf{S}_{\text{inc}}\|_{\text{v}} w_{\mathbf{r}}(\mathbf{s}_j) + \mu_{\text{a}} \|\mathbf{S}_{\text{inc}}\|_{\text{v}} w_{\mathbf{r}}(\mathbf{s}_j) + \mathbf{s}_j \cdot \nabla \|\mathbf{S}_{\text{inc}}\|_{\text{v}} w_{\mathbf{r}}(\mathbf{s}_j) \\ + \mu_{\text{s}} \|\mathbf{S}_{\text{inc}}\|_{\text{v}} w_{\mathbf{r}}(\mathbf{s}_j) - \mu_{\text{tot}} \int_{4\pi} \|\mathbf{S}_{\text{inc}}\|_{\text{v}} w_{\mathbf{r}}(\mathbf{s}') p(\mathbf{s}', \mathbf{s}_j) d\Omega'. \end{aligned} \quad (\text{A.15})$$

Equation (A.15) is the radiative transport equation for the time-averaged incoming energy flux $\langle \mathbf{S}_{\text{inc}} \rangle$ flowing in direction \mathbf{s}_j , averaged over a volume δV . It is important to note that the factor μ_{tot} from the last term appears through our definition of the phase function (2.27). It is the product $\mu_{\text{tot}} p(\mathbf{s}, \mathbf{s}_0)$ that yields the contribution which is solely due to scattering, i.e. the absorption term related to the phase function that appears through $\mu_{\text{tot}} = \mu_{\text{s}} + \mu_{\text{a}}$ has no direct physical meaning.

By defining the specific intensity as

$$I(\mathbf{r}, t, \mathbf{s}) = \frac{1}{4\pi} \|\mathbf{S}_{\text{inc}}\|_{\text{v}} w_{\mathbf{r}}(\mathbf{s}), \quad (\text{A.16})$$

equation (A.15) recovers its well-known appearance

$$\frac{1}{v} \frac{\partial I(\mathbf{r}, t, \mathbf{s})}{\partial t} + \mathbf{s} \cdot \nabla I(\mathbf{r}, t, \mathbf{s}) - (\mu_s + \mu_a) I(\mathbf{r}, t, \mathbf{s}) - \mu_{\text{tot}} \int_{4\pi} I(\mathbf{r}, t, \mathbf{s}') p(\mathbf{s}, \mathbf{s}') d\Omega' = 0. \quad (\text{A.17})$$

It is convenient to summarize the main approximations that we applied in order to reach equation (A.15). Namely they are

- average of the energy flow: $\langle \mathbf{S}_{\text{inc}}(\mathbf{r}) \rangle \simeq \|\mathbf{S}_{\text{inc}}(\mathbf{r})\|_v = \frac{1}{\delta V} \int_{\delta V} S(\mathbf{r} - \mathbf{r}') d^3 r'$
- average of the incoming energy flow much larger than the local average scattered flow much greater than local average scattered flow: $\frac{1}{\delta V} \int_{\delta V} \langle \mathbf{S}_{\text{inc}} \rangle \cdot \mathbf{s} dV \gg \frac{1}{\delta V} \int_{\delta V} \langle \mathbf{S}_{\text{sc}} \rangle \cdot \mathbf{s} dV$
- average of the gradient of the energy flow: $\|\nabla \mathbf{S}_{\text{inc}} \cdot \mathbf{s}_j\| \simeq \nabla \|\mathbf{S}_{\text{inc}}\|_v$, which directly implies that $\int_{\delta V} \nabla_{\mathbf{r}'} S(\mathbf{r} - \mathbf{r}') d^3 r' \simeq \nabla_{\mathbf{r}} \int_{\delta V} S(\mathbf{r} - \mathbf{r}') d^3 r'$
- incoherent scattering: $\langle \mathbf{S}_{\text{sc}} \rangle \simeq \sum_{i=1}^N \langle \mathbf{S}_{\text{sc}} \rangle_i$
- statistically equivalent properties throughout the medium: $\mu_{a,s} = \frac{N}{\delta V} \sigma_{a,s} \simeq \frac{N_{\text{tot}}}{V} \sigma_{a,s}$
- far-field approximation: $\langle \mathbf{S}_{\text{sc}} \rangle_i \simeq \langle \mathbf{S}_{\text{inc}}(\mathbf{r}_i) \rangle \frac{p(\mathbf{s}_i, \mathbf{s})}{|\mathbf{r} - \mathbf{r}_i|^2} d\mathbf{s}$ (see equation (2.21))
- neglect depolarization: $\nabla(\nabla \cdot \mathbf{E}) \simeq 0$

Together, these assumptions give an idea of the required size of the averaging volume δV . In order to satisfy the conditions that we set throughout this derivation, δV should be a volume much smaller the total system, but still sufficiently large to contain a large number of particles in order to be characterized by statistically meaningful optical properties.

A.2 Derivation of the diffusion equation

In this appendix we provide more details on the derivation of the diffusion equation following the P_1 approximation. In order to obtain the expansion (2.72) from the guess (2.71), we have to determine the expressions for the functions f_0 and f_1 . The isotropic term can be obtained by plugging equation (2.71) into our definition of $U(\mathbf{r}, t)$

$$U(\mathbf{r}, t) = f_0(\mathbf{r}, t) \int_{\Omega} d\Omega + f_1(\mathbf{r}, t) \int_{\Omega} \mathbf{s}_j \cdot \mathbf{s} d\Omega. \quad (\text{A.18})$$

Using $\int_{\Omega} d\Omega = 4\pi$ and $\int_{\Omega} \mathbf{s}_j \cdot \mathbf{s} d\Omega = 2\pi \int_{-1}^1 \cos \theta d\cos \theta = 0$ we obtain $f_0(\mathbf{r}, t) = U(\mathbf{r}, t)$ as expected. Analogously, we can use our ansatz in the definition of $\mathbf{F}(\mathbf{r}, t)$, which we assumed to be pointed in direction \mathbf{s}_j , to obtain

$$\mathbf{F}(\mathbf{r}, t) \cdot \mathbf{s}_j = f_0(\mathbf{r}, t) \int_{\Omega} \mathbf{s}_j \cdot \mathbf{s} d\Omega + f_1(\mathbf{r}, t) \int_{\Omega} (\mathbf{s}_j \cdot \mathbf{s})^2 d\Omega = 2\pi \int_{-1}^1 \cos^2 \theta d\cos \theta = \frac{4\pi}{3} \quad (\text{A.19})$$

and therefore $f_1(\mathbf{r}, t) = \mathbf{F}(\mathbf{r}, t)4\pi/3$.

The second approximation leading to the diffusive equation concerns the derivation of the time-dependent Fick's law, which is obtained by multiplying the RTE by \mathbf{s} and integrating over the whole solid angle. By doing so, we obtain the following terms

$$\begin{aligned} & \frac{1}{v} \frac{\partial}{\partial t} \mathbf{F}(\mathbf{r}, t) + \int_{\Omega} \left(\mathbf{s} \cdot \nabla \left[\frac{1}{4\pi} U(\mathbf{r}, t) + \frac{3}{4\pi} \mathbf{F}(\mathbf{r}, t) \cdot \mathbf{s} \right] \right) \mathbf{s} d\Omega + \mu_{\text{tot}} \mathbf{F}(\mathbf{r}, t) \\ & - \mu_{\text{tot}} \int_{\Omega, \Omega'} p(\mathbf{s}, \mathbf{s}') \left[\frac{1}{4\pi} U(\mathbf{r}, t) + \frac{3}{4\pi} \mathbf{F}(\mathbf{r}, t) \cdot \mathbf{s} \right] \mathbf{s} d\Omega' d\Omega = \int_{\Omega} q(\mathbf{r}, t, \mathbf{s}) \mathbf{s} d\Omega \quad (\text{A.20}) \end{aligned}$$

where we have introduced the terms of the P_1 expansion (2.72). In the second term of this expression we have two integrals of the form

$$\int_{\Omega} (\mathbf{s} \cdot \mathbf{A}) \mathbf{s} d\Omega = \frac{4\pi}{3} \mathbf{A} \quad (\text{A.21})$$

$$\int_{\Omega} [\mathbf{s} \cdot \nabla (\mathbf{A} \cdot \mathbf{s})] \mathbf{s} d\Omega = 0 \quad (\text{A.22})$$

which are valid for any \mathbf{A} independent of \mathbf{s} . This whole term therefore reduces to $\nabla U(\mathbf{r}, t)/3$.

Regarding the fourth term, it is the integral in $d\Omega$ of

$$\begin{aligned} & \frac{\mu_s}{4\pi} U(\mathbf{r}, t) + \mu_{\text{tot}} \frac{3}{4\pi} \int_{\Omega'} p(\mathbf{s}, \mathbf{s}') \mathbf{F}(\mathbf{r}, t) \cdot \mathbf{s}' d\Omega' \\ & = \frac{\mu_s}{4\pi} U(\mathbf{r}, t) + \mu_{\text{tot}} \frac{3}{4\pi} \mathbf{F}(\mathbf{r}, t) \cdot \int_{\Omega'} p(\mathbf{s}, \mathbf{s}') \mathbf{s}_j \cdot \mathbf{s}' d\Omega'. \quad (\text{A.23}) \end{aligned}$$

where in the last term we would need the product $\mathbf{s} \cdot \mathbf{s}'$ rather than $\mathbf{s}_j \cdot \mathbf{s}'$ to recover the first moment g of the phase function. We can reach this expression by recasting $\mathbf{s}_j \cdot \mathbf{s}'$ in terms of $\mathbf{s}_j \cdot \mathbf{s}$

$$\mathbf{s}_j \cdot \mathbf{s}' = (\mathbf{s} \cdot \mathbf{s}')(\mathbf{s}_j \cdot \mathbf{s}) + \sqrt{1 - (\mathbf{s} \cdot \mathbf{s}')^2} \sqrt{1 - (\mathbf{s}_j \cdot \mathbf{s})^2} \quad (\text{A.24})$$

which eventually yields

$$\int_{\Omega'} p(\mathbf{s} \cdot \mathbf{s}') \mathbf{s}_j \cdot \mathbf{s}' d\Omega' = (\mathbf{s}_j \cdot \mathbf{s}) \int_{\Omega'} p(\mathbf{s} \cdot \mathbf{s}') \mathbf{s} \cdot \mathbf{s}' d\Omega' = (\mathbf{s}_j \cdot \mathbf{s}) \frac{\mu_s}{\mu_{\text{tot}}} g. \quad (\text{A.25})$$

The fourth term becomes then

$$\int_{\Omega} \left[\frac{\mu_s}{4\pi} U(\mathbf{r}, t) + \frac{3g\mu_s}{4\pi} \mathbf{F}(\mathbf{r}, t) \cdot \mathbf{s} \right] d\Omega = \mu_s g \mathbf{F}(\mathbf{r}, t). \quad (\text{A.26})$$

Putting all the pieces together we finally obtain

$$\frac{1}{v} \frac{\partial}{\partial t} \mathbf{F}(\mathbf{r}, t) + \frac{\nabla U(\mathbf{r}, t)}{3} + (\mu_{\text{tot}} - g\mu_s) \mathbf{F}(\mathbf{r}, t) = \int_{\Omega} q(\mathbf{r}, t, \mathbf{s}) \mathbf{s} d\Omega \quad (\text{A.27})$$




Fig. A.1.: (a) Schematic of the boundary between a scattering and a non scattering region. Snell's law relates the incident and reflected angles at the interface. (b) Radiometric quantities such as $\mathbf{F}(\mathbf{r}, t)$ are conveniently decomposed along the tangential and normal direction to the boundary. The decomposition has a cylindrical symmetry around ϕ . (c) Coefficient $A(n)$ plotted for different values of the refractive index contrast.

At this stage, we introduce the second main assumption, namely that the temporal variation of the flux is negligible with respect to the vector itself

$$\frac{1}{\mu'_s v} \left| \frac{\partial \mathbf{F}(\mathbf{r}, t)}{\partial t} \right| \ll |\mathbf{F}(\mathbf{r}, t)|. \quad (\text{A.28})$$

It is worth commenting further this approximation, since by removing the temporal dependence of $\mathbf{F}(\mathbf{r}, t)$ we are effectively invalidating one of the fundamental similarity relations of the RTE, relating the specific intensity in the presence of absorption with the specific intensity in a non-absorbing medium (2.53). Few authors point out that the diffusion approximation is expected to hold for light that has undergone a multitude of scattering events, and therefore absorption frustrates the diffusive regime in that it selectively extinguishes the light that could propagate into the diffusive regime. In a steady state detection scheme, the acquired signal would be dominated by low-order scattering and be poorly modeled by the diffusive approximation [7, 59].

By doing so, and assuming an isotropic source $q(\mathbf{r}, t, \mathbf{s}) = q(\mathbf{r}, t)/4\pi$, Fick's law is obtained as

$$\mathbf{F}(\mathbf{r}, t) = -\frac{1}{\mu_{\text{tot}} - g\mu_s} \frac{\nabla U(\mathbf{r}, t)}{3} = -\frac{1}{v} D_{P1} \nabla U(\mathbf{r}, t) \quad (\text{A.29})$$

with $D_{P1} = \frac{v}{3(\mu_{\text{tot}} - g\mu_s)} = \frac{v}{3(\mu_a + \mu'_s)}$ as the diffusion coefficient in the P_1 approximation.

A.3 Boundary conditions at the interface of a scattering medium

The radiative transfer boundary condition between a scattering and a non scattering material is summarized by equation (2.92)

$$-\int_{\mathbf{s} \cdot \mathbf{q} < 0} I(\mathbf{r}, t, \mathbf{s})(\mathbf{s} \cdot \mathbf{q}) d\Omega = \int_{\mathbf{s} \cdot \mathbf{q} > 0} R(\theta_i) I(\mathbf{r}, t, \mathbf{s})(\mathbf{s} \cdot \mathbf{q}) d\Omega \quad (\text{A.30})$$

where $I(\mathbf{r}, t, \mathbf{s})$ is the specific intensity at the boundary Σ of the diffusive medium, $R(\theta_i)$ is the Fresnel reflection coefficient for unpolarized light and \mathbf{q} is the unit vector normal to Σ in \mathbf{r} (see Figure A.1a). In order to solve the integrals in equation (A.30) we substitute the P_1 expansion of the specific intensity (2.72) and decompose the flux as $\mathbf{F}(\mathbf{r}, t) = F_u \mathbf{u} + F_q \mathbf{q}$ (Figure A.1b), where \mathbf{u} is a unit vector tangential to Σ and $\mathbf{s} \cdot \mathbf{q} = \cos \theta_i$ and $\mathbf{u} \cdot \mathbf{s} = \sin \theta_i \cos \phi$, obtaining respectively

$$\begin{aligned} \int_{\mathbf{s} \cdot \mathbf{q} < 0} I(\mathbf{r}, t, \mathbf{s})(\mathbf{s} \cdot \mathbf{q}) d\Omega &= \frac{1}{4\pi} \int_{\mathbf{s} \cdot \mathbf{q} < 0} [U + 3\mathbf{F} \cdot \mathbf{s}](\mathbf{s} \cdot \mathbf{q}) d\Omega \\ &= \frac{U}{4\pi} \int_{\pi/2}^{\pi} \cos \theta_i \sin \theta_i d\theta_i + \frac{3}{4\pi} F_u \int_{2\pi} (\mathbf{u} \cdot \mathbf{s})(\mathbf{s} \cdot \mathbf{q}) d\Omega + \frac{3}{4\pi} F_q \int_{2\pi} (\mathbf{s} \cdot \mathbf{q})^2 d\Omega \\ &= -\frac{U}{4} + \frac{3}{4\pi} F_u \int_0^{2\pi} d\phi \int_{\pi/2}^{\pi} \cos \phi \cos \theta_i \sin^2 \theta_i d\theta_i + \frac{3}{4\pi} F_q \int_0^{2\pi} d\phi \int_{\pi/2}^{\pi} \cos^2 \theta_i \sin \theta_i d\theta_i \\ &= -\frac{U}{4} + \frac{3}{4\pi} \left(F_q \frac{2\pi}{3} \right) \end{aligned} \quad (\text{A.31})$$

and

$$\begin{aligned} \int_{\mathbf{s} \cdot \mathbf{q} > 0} R(\theta_i) I(\mathbf{r}, t, \mathbf{s})(\mathbf{s} \cdot \mathbf{q}) d\Omega &= \int_{\mathbf{s} \cdot \mathbf{q} > 0} R(\theta_i) [U + 3\mathbf{F} \cdot \mathbf{s}](\mathbf{s} \cdot \mathbf{q}) d\Omega \\ &= \frac{U}{2} \int_0^{\pi/2} R(\theta_i) \cos \theta_i \sin \theta_i d\theta_i + \frac{3}{4\pi} F_q \int_0^{2\pi} d\phi \int_0^{\pi/2} R(\theta_i) (\mathbf{s} \cdot \mathbf{q})^2 d\Omega \\ &= \frac{U}{2} \int_0^{\pi/2} R(\theta_i) \cos \theta_i \sin \theta_i d\theta_i + \frac{3}{2} F_q \int_0^{\pi/2} R(\theta_i) \cos^2 \theta_i \sin \theta_i d\theta_i \end{aligned} \quad (\text{A.32})$$

where we have used the fact that, for symmetry reasons

$$\int_{2\pi} R(\theta_i) (\mathbf{u} \cdot \mathbf{s})(\mathbf{s} \cdot \mathbf{q}) d\Omega = \int_0^{2\pi} d\phi \int_0^{\pi/2} R(\theta_i) \cos \phi \sin^2 \theta_i \cos \theta_i d\theta_i = 0. \quad (\text{A.33})$$

Substituting the expressions obtained in (A.31) and (A.32) into (A.30) and denoting the coefficient A as

$$A = \frac{1 + 3 \int_0^{\pi/2} R(\theta_i) \cos^2 \theta_i \sin \theta_i d\theta_i}{1 - 2 \int_0^{\pi/2} R(\theta_i) \cos \theta_i \sin \theta_i d\theta_i}, \quad (\text{A.34})$$

we finally obtain the partial current boundary condition (PCBC) at the interface between a diffusive and a transparent media

$$[U(\mathbf{r}, t) - 2A\mathbf{F}(\mathbf{r}, t) \cdot \mathbf{q}]_{\mathbf{r} \in \Sigma} = 0. \quad (\text{A.35})$$

Large-scale generation of exponentially distributed random numbers

In this appendix we discuss and evaluate the effect of floating-point representation on the precision of generated random numbers, and how this negatively impacts the correct sampling of the statistics of an exponential distribution. A basic approach to substantially mitigate this effect is presented, which we have implemented in MCPLUSPLUS.

In a computer, floating-point numbers can be represented with different precision. Typical examples are the `float` and `double` representations, using respectively 32 and 64 bits. Since there is only a finite number of possible sequences of bits, only a finite list of numbers in \mathbb{R} can be actually represented. In order to cover efficiently many orders of magnitude, all standard floating-point representations are organized exponentially, with some bits reserved for the significand part and some for the exponent. This means that the ‘density’ of available representations is uniform on a logarithmic scale, i.e., with an exponentially increasing accuracy towards 0.

In Monte Carlo methods, as discussed in subsection 2.1.6, a statistical distribution is sampled using uniformly distributed random numbers $\xi \in [0, 1)$, exploiting the definition of the inverse cumulative distribution. These random numbers are built as follows: the raw stream of bits obtained from the PRNG is interpreted in groups of 32 to form positive numbers `uint`, i.e., all integer numbers ζ from 0 to $n = 2^{32} - 1 = 4\,294\,967\,295$. Ideally, this list of integers can be turned into a list of equispaced fractions $0 \leq \xi < 1$ as

$$\xi = \frac{\zeta}{2^{32}} \quad (\text{B.1})$$

which must be cast into a suitable floating-point representation. However, as discussed before, fractions closer to 0 will be more precisely represented than fractions close to 1 and in general, the cast operation will involve some rounding of the exact value of the fraction to the closest available floating-point representation. As an example, let us consider the smallest and highest fractions that can be ideally defined

$$\begin{aligned} \xi_0 &= \frac{0}{2^{32}} = 0 \\ \xi_1 &= \frac{1}{2^{32}} \sim 0.000\,000\,000\,232\,830\,643\,653\,869\,629\dots \\ \xi_2 &= \frac{2}{2^{32}} \sim 0.000\,000\,000\,465\,661\,287\,307\,739\,258\dots \\ &\vdots \\ \xi_{n-1} &= \frac{2^{32}-2}{2^{32}} \sim 0.999\,999\,999\,534\,338\,712\,692\,260\,742\dots \end{aligned}$$

$$\xi_n = \frac{2^{32} - 1}{2^{32}} \sim 0.999\,999\,999\,767\,169\,356\,346\,130\,371\dots$$

A problem becomes already apparent at this stage. Even if we were able to represent these fractions with infinite precision, the fact that we cannot generate a random number arbitrarily close to 1 introduces a truncation in all derived distributions, and especially in the exponential case. If we recall equation (2.61)

$$\ell = -l_s \ln(1 - \xi) \quad (\text{B.2})$$

and substitute $\xi = \xi_n$ we obtain the maximum step length that can be drawn as $\ell_{\max} \sim -l_s \ln(1 - \xi_n) \sim 22.180\,71 \times l_s$, at which a truncation occurs. This can be easily understood if we imagine having a 2^{32} -faced die. If we want the die to be fair, then there should be no face having a probability lower (or higher) of 2^{-32} to be drawn. In fact, 2^{-32} is exactly equal to the probability of taking a step longer than ℓ_{\max} .

The situation becomes even worse if we cannot properly represent ξ_n with sufficient precision. This is the case, for example, when we cast the fractions to a 32-bit float precision, as it is sometimes done in GPU software to take advantage of their optimal performance. While the lower non-null fraction can be represented exactly as

$$\begin{array}{c} \text{sign} \quad \text{exponent} \quad \text{mantissa} \\ \boxed{0} \boxed{01011111} \boxed{000000000000000000000000} = \frac{1}{2^{32}} \end{array}$$

the largest possible representation < 1 is given by

$$\begin{array}{c} \text{sign} \quad \text{exponent} \quad \text{mantissa} \\ \boxed{0} \boxed{01011110} \boxed{111111111111111111111111} = (2^0 + 2^{-1} + \dots + 2^{-22} + 2^{-23}) \times 2^{-1} \end{array}$$

which is actually equal to 0.999 999 940 395, rather than 0.999 999 999 767. If we plug this value into equation (B.2), we actually obtain an even smaller ℓ_{\max} of just $\sim 16.6 \times l_s$, corresponding to a probability of 2^{-24} . This means that the effects of truncation start to be appreciable when more than $\sim 10^7$ steps are drawn. Conversely, when casting to 64-bit double, all fractions ξ_i can be represented accurately and the nominal truncation $\ell \leq 22.18 \times l_s$ holds. It should be noted that in double representation, numbers much closer to 1 than 0.999 999 999 767 can in principle be *represented* (double notation offers approximately 16 significant digits). The problem is that they cannot be *drawn* uniformly between $[0, 1)$ if only 2^{32} values are to be used.

Therefore, the effect of truncation comes from the combination of two effects: one is the coarse discretization of the $[0, 1)$ interval imposed by a 32-bit PRNG, and the other is the (possible) lack of accurate internal representations of these floating-point values. Several strategies are possible to lift or at least mitigate this truncation effect. In the following, we describe one straightforward and general method to increase significantly the accuracy of the simulation, which we adopted in MCPLUSPLUS. However other approximate techniques could in principle be implemented specifically for the exponential distribution, exploiting its characteristic memorylessness. The first correction consist of switching from a 32 to a 64-bit PRNG, meaning that the random bit stream generated

by the PRNG will be interpreted in groups of 64, i.e., as a positive ulong integer $\zeta = [0, 2^{64} - 1]$. This will of course halve the PRNG output rate which, however, impacts negligibly on the overall simulation time. Now the closest fraction to 1 becomes

$$\xi_n = \frac{2^{64} - 1}{2^{64}} \sim 0.999\,999\,999\,999\,999\,999\,999\,999\,945\,789\,891 \quad (\text{B.3})$$

leading ideally to a doubled step-length truncation of $\ell_{\max} \sim 44.36 \times l_s$. Step lengths longer than this would occur with a probability of $\sim 10^{-20}$, allowing to draw safely a much larger number of steps, and therefore to simulate more energy packets. As in the previous case, however, care has to be taken to check if this fraction can be represented accurately. In the common double notation, this is not the case: the closest double number preceding 1 is given by

$$\begin{array}{c} \text{sign} \quad \text{exponent} \quad \text{mantissa} \\ \boxed{0} \boxed{01111111110} \boxed{111111 \cdots 111111} = (2^0 + 2^{-1} + \cdots + 2^{-51} + 2^{-52}) \times 2^{-1} \end{array}$$

which amounts to just $\sim 0.999\,999\,999\,999\,999\,999\,999\,999\,888\,977\,697\,537 \rightarrow \ell_{\max} \sim 36.7 \times l_s$. Following the analogy with the previous case, to avoid this misrepresentation, both the equispaced fractions obtained by the 64-bit PRNGs and the logarithm of equation (B.2) can be temporarily cast to 128-bit long double precision where more than 34 significant digits are available. As can be imagined, arithmetic operations performed with 128-bit precision are much more expensive even on CPUs. However, after that the step lengths have been drawn, their values and in general the whole simulation can be normally run in double precision, since the only critical operation is the passage from the uniform ξ to the exponentially distributed ℓ .

Acknowledgments

” *The cure for boredom is curiosity. There is no cure for curiosity.*

— E. Parr

The years of my PhD have been extraordinarily rich and stimulating, and whatever the future holds, I know I will miss this heady feeling sooner or later. My gratitude goes to all the people that actually made it possible.

Diederik, your unconditioned encouragement provided me with constant motivation throughout all these years. Your work and attitude have been a great source of inspiration, but above all, I want to thank you for having let me pursue any research interest, satellite idea or side project in complete freedom. This is ultimately the reason why I feel so positively engaged with all the things that I've studied. So much so, that I want to take this opportunity to acknowledge all the incredibly talented people and groups around the world who have advanced this research field through their brilliant ideas and experiments. If this thesis has any qualities, they are largely inherited from their work. Being able to read and fully appreciate it provided me with so many moments of true amazement and inspiration than I can ever express.

As a member of Diederik's group, I had the privilege to meet and interact directly with few of these people. Starting with Matteo, who spent a great deal of his time and resources to provide us with a reference point during the first part of my PhD. Working with you has been the most formative part of all these years. Thanks for all the enthusiasm, challenges and knowledge that you threw at us. Another special mention goes to Romolo, who actually introduced me to the time-resolved setup that I kept developing. I've always admired your sensible approach to physics, and enjoyed all the good time that we had together.

Yet, it's all the other members of the group that made each day of this PhD so nice and enjoyable. Thanks to Sara, who shared with me (till the very last evening!) this rough PhD course, but would always greet you with a warm smile every morning. To Lobre, Fra and Sgrigno, the best conference companions ever, for all the fun, science and *pastis* that we had together. Thanks to ever-resourceful Dmitry, who carries on the endless fight against the setup, for your nipping sense of humor and storytelling flair. To Anjani, Hao, Filippo, Sepideh and Hua for your good company and collaboration, and to Simone, for his proactiveness and the vast experience he shared with us. Last but not least, thanks to Daniele and Camilla for being constantly available, and for all the things you take care of. Working with you is what makes me most feel part of a 'group'. Finally, I want to thank our visiting guests Jeroen, for his valuable help in the lab, and Amos, for

his many scientific and human qualities. I hope there will be many more opportunities to spend time and work together.

A special thank goes to Giacomo, a longtime friend who shared with me a lot more than just these years and his multiple talents. Thank you for all the time that we wasted together, and for all the exciting things that came out from there. I wish to extend my thanks to all his former group, and especially to Costanza, for making me feel always welcome and for her continuous support and supervision during the ups and downs of our research. Thanks to Andrea, Sofia and Pietro, for all the passions and interests that you cultivate, spanning well beyond physics. I'm glad for all the engaging time and conversations we had together.

Many more people than I could mention here, however, deserve proper acknowledgement for the time, help and resources that they gave me. I want to thank in particular Dr. Jacopo Bertolotti for hosting me in Exeter, and for all the things that I was able to learn and apply myself while I was there, and Erik Alerstam and Tomas Svensson for all the fruitful information and advice I would always get from them.

Finally, let me thank all the people that accompanied these years of study, and have become dear friends. The best part of physics is truly the wacky, passionate and compelling gang of people who do it. Not least, thanks to all my friends outside here, I wouldn't have made through all this without you. Thanks to my family, who always granted me the best conditions to do whatever I cared for, whenever I wanted, and finally, to Eva, for how it feels to be with you.

The European Research Council is gratefully acknowledged for its continuous financial support through the Advanced Grant PhotBots (proj. ref. 291349).

Colophon

This thesis was typeset with \LaTeX , using Bitstream Charter/ \TeX Gyre Heros font faces and a customized version of the *Clean Thesis* style developed by Ricardo Langner. Graphics and illustrations have been created using the packages *PGF/Tikz* by Till Tantau and *PGFplots* by Christian Feuersänger, the *colorbrewer* color set developed by Cynthia Brewer (implementation by Vincent Traag), and the perceptually uniform colormap *viridis* designed by Stéfan van der Walt and Nathaniel Smith.

