

Physics and applications of electron-matter interaction simulations

van Kessel, L.C.P.M.

DOI

[10.4233/uuid:a8c16a82-e5aa-492a-88d1-72276b8328f1](https://doi.org/10.4233/uuid:a8c16a82-e5aa-492a-88d1-72276b8328f1)

Publication date

2022

Document Version

Final published version

Citation (APA)

van Kessel, L. C. P. M. (2022). *Physics and applications of electron-matter interaction simulations*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:a8c16a82-e5aa-492a-88d1-72276b8328f1>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Physics and applications of electron-matter interaction simulations

Physics and applications of electron-matter interaction simulations

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus, Prof.dr.ir. T.H.J.J. van der Hagen,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op woensdag 19 januari 2022 om 12:30 uur

door

Lucas Cornelis Peter Maria VAN KESSEL

Master of Science in Physics & Astronomy,
Radboud Universiteit, Nederland,
geboren te Veldhoven, Nederland.

Dit proefschrift is goedgekeurd door de promotoren.

Samenstelling promotiecommissie bestaat uit:

Rector Magnificus,	voorzitter
Prof.dr.ir. P. Kruit,	Technische Universiteit Delft, promotor
Dr. C.W. Hagen,	Technische Universiteit Delft, promotor

Onafhankelijke leden:

Prof.dr. B. Rieger,	Technische Universiteit Delft
Prof.dr.ir. H. van der Graaf,	Technische Universiteit Delft
Prof.dr.ir. S.J. van der Molen,	Rijksuniversiteit Leiden
Prof.dr. W.S.M. Werner,	Technische Universität Wien
Prof.dr. L.D.A. Siebbeles,	Technische Universiteit Delft, reservelid

Overig lid:

Dr. S.F. Wuister,	ASML Netherlands
-------------------	------------------



Printed by: Ridderprint.
Front: Simulated SEM image of an impossible cube, invented by M. C. Escher for his lithograph print *Belvedere*.
Back: The same scene under a different angle.

Copyright © 2021 by L. van Kessel

ISBN 978-94-6366-493-6

An electronic version of this dissertation is available at
<https://repository.tudelft.nl/>.

Contents

Summary	vii
Samenvatting	xi
1 Introduction	1
2 Physics review	7
2.1 Scattering theory	7
2.2 Band structure model	12
2.3 Boundary crossing	13
2.4 Scattering on the electron system	15
2.5 Secondary electron generation	32
2.6 Mott scattering	36
2.7 Phonon scattering	40
2.8 Summary	50
3 Monte Carlo simulator	53
3.1 Introduction	53
3.2 Implementation	55
3.3 Reflected electrons	64
3.4 Shape of the interaction volume	67
3.5 Experimental yield comparison	69
3.6 Sensitivity analysis	72
4 Surface effects	77
4.1 Introduction	77
4.2 Surface formalism	79
4.3 Results	88
4.4 Discussion	93

4.5	Conclusions	95
5	Influence of sidewall roughness on observed line-edge roughness	97
5.1	Introduction	98
5.2	Method	99
5.3	No vertical roughness	103
5.4	Isolated lines	105
5.5	Dense lines and spaces	109
5.6	Contouring threshold	111
5.7	Energy dependence	114
5.8	Top roughness	114
5.9	Conclusions	116
6	What is missing in our simulations	119
6.1	Optical data model	120
6.2	Inner-shell excitation	125
6.3	Low-energy elastic scattering	130
6.4	Wave and band structure effects	133
6.5	Other bulk effects	134
6.6	Interface effects	137
6.7	Electron–electron interaction	140
6.8	Experimental issues	143
6.9	Discussion	146
7	Conclusion	149
	Acknowledgements	151
A	Material parameters	153
	Curriculum Vitæ	155
	Bibliography	157

Summary

Electrons with an energy ranging from 0 to 50 keV are among the most versatile tools in nanotechnology. A common example is the scanning electron microscope (SEM), which focuses an electron beam with an energy ranging from several hundred eV to tens of keV on a sample. When landing on the sample, the electrons in the beam penetrate the material. They can excite secondary electrons in the material, for example by ionization. Some of the electrons escape the sample again and reach a detector, where a high-resolution image of the sample is formed. Thanks to the small wavelength of electrons, a SEM is able to achieve single nanometre resolution while conventional optical microscopes are limited to hundreds of nanometres.

Another example is electron-beam lithography, where the electrons induce chemical reactions in a sensitive layer that is coated on top of a substrate. This layer, known as resist, locally undergoes a solubility change as a result. The soluble parts are then washed off before a subsequent step, such as etch, transfers the pattern onto the underlying substrate. Electron-beam lithography is used for low-volume manufacturing of computer chips, and the production of masks which are used in high-volume optical lithography. Extreme ultraviolet optical lithography works in a similar way. This method uses light instead of an electron beam, but the photons carry so much energy (92 eV) that they can ionize the resist. Much of the chemistry is done by electrons excited by such ionization events. A final example where electron-matter interaction is of interest, is focused electron-beam deposition. Here, secondary electrons emitted from the sample induce chemical reactions between the sample and a precursor gas which diffuses over the sample. It allows the creation of three-dimensional structures which cannot be made with other lithography techniques.

Their high spatial resolution makes SEMs suitable for use in inspection and metrology in semiconductor manufacturing. Critical dimensions on computer chips are currently less than 20 nm, and constraints on their variability are only

a few nm. When features are so small, the fact that the electrons from the SEM penetrate the sample, instead of reflecting directly, becomes important. This leads to an “interaction volume” of finite size in both the lateral and depth dimensions. The electrons emitted from the sample contain information from the entire interaction volume. The fact that the interaction volume has a finite size can be a nuisance when a very high resolution is desired, but it can also be used to observe buried features if the beam energy is chosen well. Ultimately, what information is available in an image depends on the details of interaction between the electrons and the sample. As on-chip features are becoming smaller and more three-dimensional, it is increasingly important to understand these details better.

Simulation allows insight into the details of the image formation process that cannot be obtained experimentally. It allows us to visualise the trajectories of electrons and acquire detailed knowledge about their statistics. We can then assess how certain aspects of complex samples, such as materials, heights or other three-dimensional features, translate to the image. We may also test how to set up the microscope, for example for choosing the optimal landing energy and determining acceptable spot sizes. Two aspects are important for such studies: accuracy of the physical models and computational speed. The latter is important because large parameter sweeps are often necessary before reliable conclusions can be drawn.

A good amount of research has previously been done to find good physical models. This thesis starts with an elaborate literature overview of these physical models. These models tend to agree very well when the electron energy exceeds 100 eV. At lower energies, the models are less certain. This is a problem for SEM applications, even when the electron beam has much more energy. Information about the sample shape is contained in secondary electrons, which have less than 50 eV of energy. We will put special thought into the model’s behaviour at low energy.

The next step is to develop a fast simulator with our best selection of physical models. It needs to be fast and provide insight into the detailed scattering processes when we want it. The simulations can be sped up by running them on a graphics processing unit (GPU). Unfortunately, GPUs can be somewhat cumbersome if we need to extract arbitrary information from the scattering process. The simulator was therefore implemented in a hybrid form: it can run on the central processing unit (CPU) or the GPU. The slower CPU version allows for more flexibility than the GPU simulator.

With a simulator in place, the next step is to investigate the influence of sur-

face plasmons on SEM images. It has experimentally been established that surface plasmons can contribute significantly to secondary electron emission. We have implemented a model for this in our simulator to assess their influence on top-down SEM images in a semiconductor processing setting. We find that inclusion of this effect is important if one is interested in energy spectra, but its influence on SEM images is very small.

It is then time to apply the simulation to a specific use-case in lithography: measurement of line-edge roughness (LER). LER is usually measured from top-down images. The presence of any roughness in the vertical direction is usually ignored. Developed photoresist features can be rough in the vertical direction. The question then arises how the full sidewall roughness translates to line-edge roughness upon measurement. Considering an isolated line feature, we find that the edge position observed by SEM is given by the outermost extrusion of the rough sidewall along the vertical direction. This “vertical averaging effect” means that the measured LER is lower than that of any slice along the sidewall. For dense patterns of lines and spaces, the SEM is less sensitive to lower layers closer to the substrate. The averaging effect is then less severe, and the measured LER is closer to that of the top layers in the line.

Finally, we will analyse the physical effects that are absent from the simulator. There is still plenty of room for model improvement, but there are practical issues as well, especially when modelling SEM images. The experimental circumstances in SEM are rarely perfectly controlled, and samples can have oxidation layers of unknown thickness, or are contaminated by the SEM itself. All of these effects can have a strong effect on the model outcomes, of a similar magnitude as the uncertainty in modelling itself.

These limitations on both the models and experiments mean that we must accept some uncertainty about quantitative results. That does not mean that the simulator is useless. Trends are often predicted correctly and we can often reliably draw qualitative conclusions, for example by repeating a study with a different set of materials. We will find that most of the conclusions in this thesis are not sensitive to that. That gives confidence that our results are generic.

Samenvatting

Elektronen met een energie tussen de 0 en 50 keV behoren tot de meest veelzijdige gereedschappen in de nanotechnologie. Een veelvoorkomend voorbeeld is de raster elektronenmicroscop (SEM), die een elektronenbundel met een energie tussen enkele honderden eV en tientallen keV op een preparaat focuseert. Wanneer de elektronen het preparaat raken, dringen ze door in het materiaal. Daar kunnen ze secundaire elektronen uit het materiaal aanslaan, bijvoorbeeld door ionisatie. Sommige elektronen kunnen het preparaat weer ontsnappen en een detector bereiken, waar met hoge resolutie een beeld van het preparaat wordt gevormd. Dankzij de kleine golflengte van elektronen kan een SEM een resolutie van enkele nanometers bereiken, terwijl conventionele optische microscopen beperkt zijn tot honderden nanometers.

Nog een voorbeeld is elektronbundellithografie, waarbij de elektronen chemische reacties induceren in een gevoelige laag die op een substraat wordt aangebracht. Deze laag, die resist wordt genoemd, ondergaat daardoor lokaal een oplosbaarheidsverandering. De oplosbare delen worden daarna weggespoeld voordat een volgende stap, zoals ets, het patroon overbrengt naar het onderliggende substraat. Elektronbundellithografie wordt gebruikt voor productie van computerchips in kleine volumes, en voor de productie van maskers die gebruikt worden in optische lithografie voor hoge volumes. Extreem ultraviolette optische lithografie werkt op een soortgelijke manier. Bij deze methode wordt licht gebruikt in plaats van een elektronenbundel, maar de fotonen dragen zoveel energie (92 eV) dat ze de resist kunnen ioniseren. Een groot deel van de chemie wordt gedaan door de elektronen die door deze ionisatie worden aangeslagen. Een laatste voorbeeld waarbij de interactie tussen elektronen en materie interessant is, is elektrongeïnduceerde depositie. Hierbij induceren secundaire elektronen die uit een preparaat komen chemische reacties tussen het preparaat en een precursorgas die over het preparaat diffundeert. Hiermee kunnen driedimensionale structuren worden gemaakt die niet met andere li-

thografische technieken kunnen worden geproduceerd.

Hun hoge ruimtelijke resolutie maakt SEMs nuttig voor inspectie en metrologie in halfgeleiderfabricage. Kritieke dimensies op computerchips zijn momenteel kleiner dan 20 nm, en de maximaal toegestane variabiliteit is slechts enkele nm. Als de structuren zo klein zijn, is het feit dat de elektronen van de SEM doordringen in het preparaat en niet direct reflecteren belangrijk. Dit leidt tot een “interactievolume” van eindige grootte in zowel de laterale als diepterichting. De elektronen die uit het preparaat komen bevatten informatie uit het gehele interactievolume. Het feit dat het interactievolume een eindige grootte heeft kan vervelend zijn als een zeer hoge resolutie gewenst is, maar het kan ook nuttig zijn om begraven structuren te inspecteren indien de bundelenergie goed gekozen wordt. Uiteindelijk hangt de informatie die het beeld bevat af van de gedetailleerde interactie tussen de elektronen en het preparaat. Aangezien de structuren op chips kleiner en driedimensionaler worden, wordt het ook steeds belangrijker om deze details beter te begrijpen.

Simulatie geeft inzicht in de details van het beeldvormingsproces die experimenteel niet verkregen kunnen worden. Het staat ons toe om trajecten van elektronen te visualiseren en gedetailleerde kennis te verwerven over hun statistiek. We kunnen dan beoordelen hoe bepaalde aspecten van complexe preparaten, zoals materiaal, hoogte of driedimensionale structuren, zich vertalen naar een beeld. We kunnen ook toetsen hoe we de microscoop moeten instellen, bijvoorbeeld om de optimale bundelenergie te kiezen en acceptabele bundelgroottes te bepalen. Twee aspecten zijn belangrijk voor dergelijke studies: accuratesse van de fysische modellen en rekensnelheid. De laatste is belangrijk omdat het vaak nodig is om een grote parameter ruimte te bestrijken voordat er betrouwbare conclusies getrokken kunnen worden.

Er is in het verleden aardig wat onderzoek gedaan naar goede fysische modellen. Dit proefschrift begint met een uitgebreid literatuuroverzicht over deze fysische modellen. Deze modellen komen vaak zeer goed overeen als de energie van de elektronen boven 100 eV zit. Bij lagere energieën zijn de modellen minder zeker. Dit is een probleem voor SEM-toepassingen, zelfs als de bundel zelf veel meer energie bezit. De informatie over de vorm van het preparaat zit in secundaire elektronen, die minder dan 50 eV aan energie hebben. We zullen speciale aandacht besteden aan het gedrag van onze modellen bij deze lage energieën.

De volgende stap is het ontwikkelen van een snelle simulator met onze beste selectie van fysische modellen. Deze moet snel zijn en inzicht verschaffen in de gedetailleerde verstrooiingsprocessen wanneer we die willen. De simulaties

kunnen worden versneld door ze op een videokaart (GPU) te draaien. Helaas kan het gebruik van GPUs nogal omslachtig zijn als we arbitraire informatie over het verstrooiingsproces moeten meten. De simulator is daarom in een hybride vorm geïmplementeerd: hij kan op de centrale processor (CPU) of op de GPU draaien. De langzamere CPU-versie staat meer flexibiliteit toe dan de GPU-simulator.

Als de simulator af is, is de volgende stap om de invloed van oppervlakteplasmonen op SEM-beelden te onderzoeken. Er is experimenteel vastgesteld dat oppervlakteplasmonen significant kunnen bijdragen aan de emissie van secundaire elektronen. We hebben hier een model voor in onze simulator geïmplementeerd om hun invloed op SEM-beelden in een halfgeleiderfabricageomgeving te beoordelen. We komen tot de slotsom dat het meenemen van dit effect belangrijk is als men geïnteresseerd is in energiespectra, maar dat de invloed op SEM-beelden zeer klein is.

Het is dan tijd om de simulaties toe te passen om een specifieke situatie die in de lithografie voorkomt: het meten van lijnruwheid (LER). LER wordt meestal vanuit bovenaanzicht gemeten. De aanwezigheid van ruwheid in de verticale richting wordt meestal genegeerd. Structuren in ontwikkelde fotolak kunnen ruwheid in de verticale richting vertonen. Dat werpt de vraag op hoe de volledige ruwheid van de zijmuur zich vertaalt naar lijnruwheid. Kijkend naar een losse lijn, concluderen we dat de positie die de SEM observeert wordt gegeven door het buitenste punt van de zijmuur, gezien langs de verticale richting. Als gevolg van dit middelingseffect in de verticale richting is de gemeten LER lager dan die van een willekeurig plakje van de zijmuur. In het geval van een dicht patroon van lijnen en leegtes is de SEM minder gevoelig voor de diepere lagen die dicht bij het substraat liggen. Het middelingseffect is dan minder zwaar, en de gemeten LER licht dicht bij die van de bovenste lagen van de lijn.

Ten slotte analyseren we de fysische effecten die de simulator nog niet meeneemt. Er is ruimschoots ruimte voor verbetering, maar er zijn ook praktische complicaties, vooral bij het modelleren van SEM-beelden. De experimentele omstandigheden in een SEM worden zelden goed beheerst, en preparaten kunnen oxidatielagen van onbekende dikte hebben of worden door de SEM zelf vervuild. Al deze effecten kunnen een sterke invloed op de modeluitkomsten hebben, van soortgelijke grootte als de onzekerheid in het model zelf.

Deze beperkingen aan zowel de modellen als de experimenten betekenen dat we een zekere onzekerheid moeten accepteren wat betreft de kwantitatieve resultaten. Dat betekent echter niet dat de simulator nutteloos is. Trends kunnen vaak correct worden voorspeld en we kunnen vaak betrouwbare kwalita-

tieve conclusies trekken, bijvoorbeeld door een onderzoek te herhalen met een andere materiaalkeuze. We zullen zien dat de meeste conclusies in dit proefschrift daar niet afhankelijk van zijn. Dat geeft vertrouwen dat de resultaten algemeen geldig zijn.

Chapter 1

Introduction

The study of electricity and magnetism has a long and storied history. Various electric and magnetic phenomena, including static electricity, magnetic minerals and lightning, have been studied for millennia. Fundamental scientific understanding of electricity and magnetism grew considerably during the eighteenth and nineteenth centuries, culminating in Maxwell's *Treatise on Electricity and Magnetism* [1].

Around the same time, it was discovered that when a high voltage was applied between two electrodes in partial vacuum, the negative electrode emits radiation known as cathode rays. A number of physicists proposed that cathode rays might consist of discrete units, called electrons. J.J. Thomson is usually credited with the experimental discovery of the electron as a particle in 1897 [2] by showing that cathode rays carried an electric charge. He later received a Nobel prize for his work. The discovery of the electron paved the way for a fundamental understanding of the structure of atoms and matter, and the associated development of quantum mechanics, in the remarkable decades that lay ahead.

After the invention of the electromagnetic lens, it made sense to design a microscope that used electrons instead of light. It had been known for some time that a light microscope was unable to resolve distances smaller than those given by the Abbe diffraction limit, which is a few hundred nanometres for visible light. Following the hypothesis of De Broglie, the diffraction limit of electrons is considerably smaller, allowing for sub-nanometre resolution. The first to construct an electron microscope was Ruska in 1931. He was able to beat the diffraction limit of light a few years later, and won the Nobel prize in 1937 [3].

In 1947, Bardeen, Brattain and Shockley built the world's first transistor. The transistor had in fact already been patented some 20 years earlier, but it is likely

1 that nobody succeeded in building one. Bardeen, Brattain and Shockley would later be awarded the Nobel prize for their work. The first transistors were relatively large and power-hungry devices, but that changed with the invention of the metal-oxide-semiconductor field-effect transistor (MOSFET) by Atalla and Kahng in 1959. The invention of the MOSFET enabled integrating massive numbers of transistors on a single integrated circuit (IC) chip. By 1965, Moore noticed that the density of components on a chip at minimum cost had doubled every year since 1959 [4]. He made the bold speculation that this exponential trend would continue for another ten years. Moore's law soon transitioned from a historical observation to a target for the semiconductor industry to achieve. Remarkably, this turned out to be possible, and Moore's law has held more or less true to the present day¹. The MOSFET and Moore's law have been central to the digital revolution of recent decades. They have directly enabled the development of (portable) personal computers, smartphones and the internet, amongst many others.

One of the main driving forces behind Moore's law has been a shrinking of the features on chips, a trend which continues to this day. This shrinking is driven by continuous improvement of the lithographic techniques used for manufacturing. As on-chip components shrink, so do the constraints on the variability of their dimensions. Critical dimensions are approaching the 10 nm range, and the industry usually takes the maximum acceptable 3σ variability as 10% of that. More variability leads to reduced device performance, or potentially stochastic defects, meaning that a feature fails to print. Process control therefore becomes increasingly important. This comprises both metrology, that is, measurements of feature sizes and their variability, and inspection, to analyse specific defects. It makes sense to use electron microscopy—more specifically, the scanning electron microscopy (SEM) variety—for this purpose. There are alternative techniques which can achieve similar resolutions, such as transmission electron microscopy (TEM) and atomic force microscopy (AFM). TEM uses electrons of much higher energy than SEM and performs the measurement on the transmission side instead of in reflection as SEM does. This method is destructive as the sample needs to be sliced into thin lamellae in order to allow the electrons to penetrate through the sample. AFM is non-destructive, but very slow and suffers from resolution issues in the lateral direction due to the shape of the tip.

A SEM focuses an electron beam into a nm-sized spot on a sample. The

¹The doubling time turned out to be two years, and Moore's law is known in various other guises, including transistors per chip, transistors per dollar, performance per watt, and others.

electrons in this beam then penetrate the sample and scatter through the bulk of the sample. They are able to excite valence electrons in the material, which scatter further and may eventually escape the sample as secondary electrons (SEs). The electrons from the beam may be stopped, or turn around and escape as backscattered electrons (BSEs). With the beam focused in one position, the SEM measures the current of escaping SEs or BSEs. Scanning the beam in a grid pattern produces a greyscale image, where the value of each pixel is given by the outgoing current of electrons at that point.

Images made using secondary electrons are sensitive to the topography of the surface on which they land. When the features on the sample are large, these images are often quite straightforward to interpret. Below a certain scale, however, the finite size of the interaction volume starts to play a role in the image formation process. The features that can be made with nanolithography are now so small that this should be accounted for. Backscattered electron images are not so sensitive to the surface topography, but they do carry information about the elemental composition. They also carry depth information, and depending on the energy of the primary electron beam, one can see varying depths in the sample. This is very interesting for inspection purposes, but the question then arises how the energy should be tuned to get the most information about a depth of interest. Simulation allows insight into both these problems, as well as many others, in a way that is not possible (or much more expensive) experimentally.

This thesis will focus on making a simulator of electron-matter interaction, with a special interest in applications for SEMs. The most popular method for doing that is Monte Carlo simulation. The Monte Carlo method uses random sampling for its calculations. This is a natural fit for our simulator, as the underlying physics of electron scattering is inherently random. Our simulator will enact one specific realisation of what might happen to each electron that lands on the sample. There have been attempts to use the Boltzmann transport equation instead of the Monte Carlo method [5]. These methods are less flexible when one is interested in non-trivial geometries, and it is more difficult to incorporate arbitrary scattering effects. More importantly though, deterministic methods do not provide any information on the noise in a SEM image. The noise is an interesting quantity for our purposes, as SEM images tend to be taken with the lowest acceptable dose of electrons to minimise sample damage and maximise throughput. The noise in a Monte Carlo sampled image is representative of the actual experimental noise.

An early scientific application of the Monte Carlo method was Lord Kelvin's

1

simulation of the motion of a particle in a volume with rough edges [6]. As this was done before computers were invented, Kelvin had his assistant draw cards from a deck to generate random numbers and carry out the simulation by hand. The Monte Carlo method became more workable with the advent of computers, and it was formalised by Ulam, von Neumann and co-workers as part of the Manhattan project [7]. They coined the term Monte Carlo, after the famous casino in Monaco. It seems that the first to apply the Monte Carlo method to electron scattering were Hebbard and Wilson [8] in 1955. As time progressed, computers became more powerful and simulations became increasingly detailed. Nevertheless, performing true event-by-event simulations remained prohibitively expensive for a long time. Multiple scattering events were therefore often “condensed” into a single step. Many models employed the continuous slowing-down approximation, which makes an electron gradually lose energy along its trajectory as opposed to in discrete events as is the case in reality. The plural-scattering model, in which multiple elastic events are merged into one, was also popular into the 1990s. Not until the late 2000s did computers become fast enough that per-event simulations became commonplace. In addition to merging multiple events into one, approximate models for the events themselves were often employed when more accurate but expensive alternatives were available. Many of those approximate models featured fudge factors which could be tuned to calibrate simulation output to experimental data.

A wide variety of simulators has recently appeared in literature [9–17], but such software is often not publicly or freely available, features outdated or approximate physical models, or is too slow for our needs. We will make a simulator that is fully based in physical models while being as fast as possible. We wish to eliminate all fudge factors to avoid the risk that incorrectly tuned parameters compensate each other, in which case the calibration to experiment may look fine while results for an application of interest may be off. The physics of high-energy ($\gtrsim 100$ eV) electron scattering is quite well known in literature, but the energy range below that is much more difficult to treat from a physics point of view. We will therefore pay special attention to the models used in this low energy range.

While this thesis is written with an application to SEM in mind, the simulator is by no means limited to that application. For example, it has also been used to study secondary electron emission in microchannel plates [18]. It has also been used to simulate part of the electron-beam lithography process, by assuming that the positions where electrons lose energy is where breaking of

chemical bonds is induced in positive-tone photoresists. A modified version has even been used to simulate the growth process of structures in electron beam induced deposition.

This thesis is organised as follows. Chapter 2 lays out the physical foundations upon which the simulator is based. The physics contained in this chapter is available in literature. Chapter 3 discusses the technical background of the simulator itself. We then investigate one specific physical phenomenon which is usually ignored in literature, namely the contribution of surface plasmons to SEM images. This will be done in chapter 4. It is then time to use the simulator to investigate an application in more detail in chapter 5. One problem with CD-SEM images of rough structures is that the roughness is inherently three-dimensional while the image has only two dimensions. The simulator allows us to investigate this relationship more closely. Finally, chapter 6 discusses the limitations associated with the simulations we do in this thesis, and when we might expect this simulator to give incorrect predictions.

Chapter 2

Physics review

This chapter reviews the physical theory of electron-matter interaction. Section 2.1 starts with a basic introduction to scattering theory. We then move on to the relevant physics for electron-matter interaction. Section 2.2 describes the simplified band structure model used in this work. Sections 2.4, 2.5, 2.6 and 2.7 describe relevant scattering processes. We finally summarize the models in section 2.8.

2.1 Scattering theory

Consider a point projectile moving towards a hard target. The target is positioned at the origin and the projectile moves parallel to the z axis. The projectile's distance to the z axis is called the impact parameter, b . The situation is illustrated in figure 2.1a.

If b is unknown, the probability that the projectile hits the target is proportional to the target's projected cross-sectional area. This cross-sectional area, also called the *scattering cross section*, is denoted with σ . The details of the interaction when the projectile reaches the target are also probabilistic. This probability distribution is captured by the *differential scattering cross section*. For example, if the projectile is deflected, the differential scattering cross section may be denoted as $d\sigma / d\Omega$, where $d\Omega$ represents the solid angle element into which the projectile is deflected. The differential scattering cross section integrates to the cross section (giving the total probability of scattering); if an interaction happens, the probability distribution is given by $\sigma^{-1} d\sigma / d\Omega$.

In contrast to what figure 2.1a implies, the target is not necessarily a hard impenetrable object. It can be anything that may cause the projectile to be de-




Figure 2.1: (a) A projectile moves towards a target situated at the origin. If the impact parameter b is unknown, the scattering probability is proportional to the target's projected surface area. (b) The projectile moves towards a thin slice of material. If the lateral position of the projectile is unknown, the scattering probability is equal to the ratio of target areas to the slice area.

flected or lose energy with a certain probability distribution. For example, the projectile may be a charged particle and the target may be a charge distribution. The concepts of (differential) scattering cross section translate to such situations in a straightforward manner. However, the scattering cross section is not necessarily equal to the geometrical cross-sectional area of the target.

In this thesis, the projectile is an electron, moving through a homogeneous material. In some cases, the targets will be the atoms in the material, but we will consider the material as a more generic collection of scattering centres. There are n of them per unit volume, each with the same scattering cross section σ . We make one more assumption, which is that the scattering centres are distributed independently of each other.

Consider an electron normally incident on a thin slice of material, with area A_{slice} and thickness Δz . The situation is illustrated in figure 2.1b. The probability that the electron scatters is equal to the ratio of the scattering cross section of the scattering centres in the slice to the area of the slice itself:

$$\frac{A_{\text{targets}}}{A_{\text{slice}}} = \frac{\sigma n A_{\text{slice}} \Delta z}{A_{\text{slice}}} = \sigma n \Delta z. \quad (2.1)$$

Now, assume that N electrons are incident on this slice. If the slice is thin, the number of electrons that do not scatter in the slice is equal to $N(1 - \sigma n \Delta z)$. For an infinitesimally thin slice, $\Delta z \rightarrow dz$,

$$\frac{dN}{dz} = -N\sigma n = -\frac{N}{\lambda}, \quad (2.2)$$

where we defined a new quantity $\lambda = 1/(\sigma n)$. Assuming that it is a constant, the differential equation can be solved:

$$N(z) = N_0 \exp(-z/\lambda). \quad (2.3)$$

This is the Lambert-Beer law. Setting $N_0 = 1$, this becomes a probability distribution for the electron to travel a certain path length without scattering:

$$p(z) = \frac{1}{\lambda} e^{-z/\lambda}. \quad (2.4)$$

The expectation value for this probability distribution is λ . This quantity is therefore called the *mean free path*.

The inverse mean free path, λ^{-1} , is a very convenient quantity because it is proportional to the scattering cross section σ . The concept of differential scattering cross section carries over to the *differential inverse mean free path*. Also, when there are multiple independent scattering processes, their differential scattering cross sections can be summed to give a total differential scattering cross section. The same holds for differential inverse mean free paths.

2.1.1 Inelastic scattering

In this thesis, we define an inelastic scattering event as an event where the projectile loses energy. This definition is slightly different from the traditional concept of an inelastic collision, where it is the *total* amount of kinetic energy that is not conserved.

An electron travelling through matter typically loses only a small fraction of its total energy. It undergoes a series of inelastic events before it is finally stopped. A very useful quantity for describing inelastic events is therefore the *stopping power*, or the amount of energy lost per unit path length travelled. The stopping power is given by

$$S = -\frac{dE}{dx} = \int_0^E d\omega \omega \frac{d\lambda^{-1}}{d\omega}, \quad (2.5)$$

where E represents the electron's kinetic energy, x the path length it travels. ω represents energy loss and λ^{-1} is the inverse mean free path. The right-hand side of this equation can be understood as follows. Remember that, if an event takes place, $P(\omega) = \lambda \frac{d\lambda^{-1}}{d\omega}$ represents the probability distribution of energy loss. Then $\int d\omega \omega P(\omega)$ gives the average amount of energy loss per event. Dividing this by λ gives the stopping power. The integral runs from zero to the energy of




Figure 2.2: Kinematics for inelastic scattering. The incoming particle's momentum is \mathbf{p} , its kinetic energy is T . It transfers momentum \mathbf{q} and energy ω , and ends up with momentum \mathbf{p}' and energy T' .

the electron, since it can't gain energy in a collision and it can't lose more than it has.

It is also possible to define a “stopping range”, the mean distance an electron can travel before it is stopped. This range would be given by

$$L_S = \int_0^{E_0} \frac{dE}{S(E)}. \quad (2.6)$$

Here, E_0 is the starting energy of the electron, and the integral runs over all lower energies E . However, as will be mentioned in section 2.4, the stopping power becomes very small at low energies. The stopping range is therefore not a very useful quantity in practice.

Before moving on, we will briefly derive some energy-momentum relationships which are useful later on. Consider an incoming particle with momentum \mathbf{p} and kinetic energy T . It loses momentum \mathbf{q} and energy ω , ending up with momentum $\mathbf{p}' = \mathbf{p} - \mathbf{q}$ and energy $T' = T - \omega$. See the diagram in figure 2.2.

The relativistic energy-momentum relation is

$$(T + mc^2)^2 = (pc)^2 + (mc^2)^2. \quad (2.7)$$

Meanwhile, from the definition that $\mathbf{q} = \mathbf{p} - \mathbf{p}'$,

$$(cq)^2 = (cp)^2 + (cp')^2 - 2c^2 pp' \cos \theta \quad (2.8)$$

$$= T(T + 2mc^2) + T'(T' + 2mc^2) - 2\sqrt{T(T + 2mc^2)}\sqrt{T'(T' + 2mc^2)} \cos \theta. \quad (2.9)$$

The minimum and maximum possible momentum transfers can be found by setting $\cos \theta = \pm 1$:

$$\frac{q_{\pm}}{2m} = \sqrt{T(1 + T/2mc^2)} \pm \sqrt{T'(1 + T'/2mc^2)}. \quad (2.10)$$




Figure 2.3: Illustration corresponding to the explanation of the transport length. See the main text for details.

2.1.2 Elastic scattering

In this thesis, an elastic scattering event is defined as an event where the projectile does not lose any energy, or only a negligible amount. The projectile changes its direction of travel, and continues with the same energy it had before the event. Specifically, we will consider two types of elastic scattering in this thesis. The first is deflection by atoms (section 2.6), in which the electron loses at most $\sim 10^{-4}$ of its own kinetic energy to the recoil of the atom (usually much less because forward scattering dominates). The other is scattering on acoustic phonons (section 2.7) in which around 10 meV is lost or gained depending on whether a phonon is emitted or absorbed. When elastic scattering is treated in a wave-optical formalism and the spatial coherence of the electron is sufficient, the scattering on a crystal is concentrated in only a few directions. This effect, diffraction, is not included in our model.

When an electron travelling through matter scatters elastically, it is usually deflected by a small amount. It takes a series of elastic events before the electron has lost its “knowledge” about the direction it was initially travelling in. The *transport length* quantifies the distance an electron travels before this has happened. To understand the transport length, consider figure 2.3. In the first step, the electron travels a distance L_0 into the z direction. It is deflected by an angle θ_1 , then travels a distance L_1 to the next interaction point. The distance travelled along the z direction in this step is equal to $L_1 \cos \theta_1$. The electron is then deflected by an angle θ_2 and travels a distance L_2 . Because the electron is also deflected in the azimuthal direction (indicated by the cone in the figure), the electron’s orientation with respect to the z axis is not simply given by $\theta_1 + \theta_2$.

If the azimuthal scattering direction is random and uniformly distributed, the average distance travelled along the z axis is given by the centre of the cone: $L_2 \cos\theta_1 \cos\theta_2$. Now, assuming that the mean free path doesn't change as the electron travels, and all scattering angles are independent and identically distributed, the average distance along z travelled in step i is given by

$$\langle \Delta z_i \rangle = \lambda \langle \cos\theta \rangle^i, \quad (2.11)$$

where $\langle \rangle$ denotes the average value. The transport length is now defined as the infinite sum

$$\begin{aligned} L_T &= \sum_{i=0}^{\infty} \langle \Delta z_i \rangle \\ &= \frac{\lambda}{1 - \langle \cos\theta \rangle}. \end{aligned} \quad (2.12)$$

In other words, the transport length is the projected distance along z before information about the original direction has been lost. In terms of the differential inverse mean free path, the transport length is given by

$$L_T^{-1} = \int d\Omega (1 - \cos\theta) \frac{d\lambda^{-1}}{d\Omega}. \quad (2.13)$$

2.2 Band structure model

Let us now briefly mention a simple model of a solid's electronic structure at zero temperature. The lowest-lying electrons fill core energy levels of the atoms in the solid. Because the atoms are close together, higher electron states merge into a continuum band, known as the valence band. The last filled level in the valence band is the Fermi energy. Above the Fermi energy, there are higher unfilled states, forming the conduction band. In the case of insulators or semiconductors, the conduction band is separated from the valence band by a band gap of forbidden states.

We use a highly simplified band structure model in this thesis. We choose the zero of energy to be the bottom of the valence band of the material that the electron is in. Due to the very complicated potential in a solid, electrons can have a complex dispersion relation between their momentum $\hbar\mathbf{k}$ and energy E . We ignore that in this thesis, assuming the dispersion relation of free space

$$E = \frac{\hbar^2 k^2}{2m}. \quad (2.14)$$




Figure 2.4: Diagram of the simple band structure model used. Energy levels are shown along the vertical axis. (a) is for metals, (b) for semiconductors and insulators. U is the inner potential, Ψ the work function, χ the electron affinity, G the band gap and V the width of the valence band.

Here, m represents the electron's mass.

Typically, the zero of energy in the vacuum, which we will call the vacuum level, lies above the Fermi energy. In a metal, the difference between the vacuum level and the Fermi energy is known as the work function Ψ . In an insulator or semiconductor, the difference between the vacuum level and the bottom of the conduction band is the electron affinity χ . Because we reference all energies with respect to the bottom of the valence band, it is convenient to define the *inner potential* U as the difference between the bottom of the valence band and the vacuum level. See figure 2.4 for an illustration. Note that we will speak of metals as if they have separate valence and conduction bands, even though they have only one band. We will use a (somewhat sloppy) language where the term “valence band” refers to the part of the band below the Fermi energy, where the valence electrons are, and the term “conduction band” refers to the unoccupied states above it.

2.3 Boundary crossing

When an electron crosses an interface between two materials, or a material and vacuum, it feels the change in inner potential. Let us consider an electron moving from a material with inner potential U_1 to a material with inner potential U_2 .




Figure 2.5: An electron crosses a boundary between materials 1 and 2, with inner potentials U_1 and U_2 , respectively. The electron's momentum parallel to the interface is conserved, but it changes in the perpendicular direction.

The inner potential of vacuum is zero.

The electron's momentum in the direction parallel to the interface is conserved. In the direction perpendicular to the interface, we assume that the electron sees a simple step potential, as in figure 2.5a, of height $\Delta U = U_2 - U_1$. This is a popular model [9–13], though more elaborate shapes have also been proposed [17].

If the electron does not have enough momentum in the perpendicular direction to jump the barrier, it must be reflected. But even when the electron has sufficient energy to cross the barrier, it still has a finite probability to be reflected. This is a popular problem in elementary quantum mechanics textbooks [19]. Defining the auxiliary variable

$$z = \sqrt{1 + \frac{\Delta U}{E \cos^2(\theta_1)}}, \quad (2.15)$$

the reflection and transmission coefficients are

$$R = \frac{(1 - z)^2}{(1 + z)^2}, \quad (2.16a)$$

$$T = \frac{4z}{(1 + z)^2}, \quad (2.16b)$$

respectively.

If the electron is transmitted, it gains or loses ΔU of energy. The electron's momentum in the direction parallel to the interface is conserved, and changes

by $\sqrt{2m\Delta U}$ in the perpendicular direction. The resulting change of direction is often called “refraction” of the electron, in analogy with light. If the electron’s energy in material 1 is E , then the deflection angles (see figure 2.5b) are related by

$$\sqrt{E} \sin\theta_1 = \sqrt{E + \Delta U} \sin\theta_2. \quad (2.17)$$

2.4 Scattering on the electron system

If a charged particle travels through a material, the valence electrons in the material rearrange (polarize) to screen its charge. If the charged particle is moving, the electric field induced by the valence electrons acts to slow the particle down. The exact response of the material can be described by the dielectric function. We use this to describe inelastic electron scattering in a material.

This section gives a detailed description of the theoretical foundations. This theory is well described in literature, but we would like to understand the key assumptions made. This section is limited to bulk materials, in other words, surface effects are neglected. The influence of surface effects is discussed in more detail in chapter 4.

2.4.1 Stopping of charged particles

The detailed interaction of charged particles with matter can be treated from a full quantum mechanical calculation [20]. The technicalities of such a calculation are quite involved. The results are equivalent to those from a semiclassical approach [21, 22], which we will showcase here. We will use the SI system of units.

Maxwell’s equations in matter are

$$\nabla \cdot \mathbf{D}(\mathbf{r}, t) = \rho_f(\mathbf{r}, t), \quad (2.18a)$$

$$\nabla \cdot \mathbf{B}(\mathbf{r}, t) = 0, \quad (2.18b)$$

$$\nabla \times \mathbf{E}(\mathbf{r}, t) = -\frac{\partial \mathbf{B}(\mathbf{r}, t)}{\partial t}, \quad (2.18c)$$

$$\nabla \times \mathbf{H}(\mathbf{r}, t) = \mathbf{J}_f(\mathbf{r}, t) + \frac{\partial \mathbf{D}(\mathbf{r}, t)}{\partial t}. \quad (2.18d)$$

The magnetic field \mathbf{B} and electric field \mathbf{E} can be written in terms of the vector

potential \mathbf{A} and scalar potential ϕ :

$$\mathbf{B}(\mathbf{r}, t) = \nabla \times \mathbf{A}(\mathbf{r}, t), \quad (2.19a)$$

$$\mathbf{E}(\mathbf{r}, t) = -\nabla\phi(\mathbf{r}, t) - \frac{\partial\mathbf{A}(\mathbf{r}, t)}{\partial t}. \quad (2.19b)$$

We now assume that the electron is moving slowly compared to the speed of light. This implies that the transverse components of the electromagnetic field can be ignored. It is then convenient to work in the Coulomb gauge and set $\nabla \cdot \mathbf{A}(\mathbf{r}, t) = 0$. This approximation reduces Maxwell's equations to the electrostatic case, even though the electric field strictly speaking varies with time. It is valid for electron energies much less than the electron rest mass, 511 keV.

We now switch to Fourier space, where all quantities are functions of wave vector \mathbf{q} and frequency ω instead of position and time. For each quantity Q , we will use the convention

$$Q(\mathbf{q}, \omega) = \frac{1}{(2\pi)^2} \int d\mathbf{r} \int dt e^{-i(\mathbf{q}\cdot\mathbf{r} - \omega t)} Q(\mathbf{r}, t), \quad (2.20a)$$

$$Q(\mathbf{r}, t) = \frac{1}{(2\pi)^2} \int d\mathbf{q} \int d\omega e^{i(\mathbf{q}\cdot\mathbf{r} - \omega t)} Q(\mathbf{q}, \omega). \quad (2.20b)$$

We also assume that the material responds linearly to the presence of the electron, *i.e.* the displacement field \mathbf{D} is related to the electric field \mathbf{E} by

$$\mathbf{D}(\mathbf{q}, \omega) = \epsilon_0 \epsilon(\mathbf{q}, \omega) \mathbf{E}(\mathbf{q}, \omega). \quad (2.21)$$

The quantity $\epsilon(\mathbf{q}, \omega)$ is the dielectric function of the medium. This assumption is formally equivalent to a first-order Born approximation. It has been suggested [23] that for the case of inelastic scattering in solids, this approximation is only valid for kinetic energies above $7E_F$, where E_F is the Fermi energy. For the electron energies dealt with in this thesis, it would be better to use the second-order Born approximation also given in ref. [23]. Interestingly, the second-order approximation is ignored in most literature [9–17], even in studies focusing entirely on the sub-100 eV energy range [24, 25]. In this thesis, we follow this literature convention and do not implement the second-order approximation. We will come back to this point in section 6.1.1.

Finally, we describe the electron classically. It is a point charge moving with velocity \mathbf{v} , so its charge distribution (without loss of generality, passing through the origin at time $t = 0$) is given by

$$\rho(\mathbf{r}, t) = -e \delta^{(3)}(\mathbf{r} - \mathbf{v}t). \quad (2.22)$$

With the assumptions above, let us compute the dielectric displacement due to a moving electron. From the Fourier transform of equation (2.18a),

$$\begin{aligned} i\mathbf{q} \cdot \mathbf{D}(\mathbf{q}, \omega) &= -\frac{e}{(2\pi)^2} \int d\mathbf{r} \int dt e^{-i(\mathbf{q}\mathbf{r}-\omega t)} \delta^{(3)}(\mathbf{r}-\mathbf{v}t) \\ &= -\frac{e}{(2\pi)^2} \int dt e^{-i(\mathbf{q}\cdot\mathbf{v}-\omega)t} \\ &= -\frac{e}{2\pi} \delta(\mathbf{q} \cdot \mathbf{v} - \omega). \end{aligned} \quad (2.23)$$

The Fourier transform of (2.19b) (with $\mathbf{A} = 0$) is

$$\mathbf{E}(\mathbf{q}, \omega) = -i\mathbf{q} \phi(\mathbf{q}, \omega). \quad (2.24)$$

Since $\mathbf{D} \parallel \mathbf{E} \parallel \mathbf{q}$, we can extract the electric field from the displacement field:

$$\mathbf{E}(\mathbf{q}, \omega) = \frac{ie}{2\pi\epsilon_0\epsilon(\mathbf{q}, \omega)} \frac{\mathbf{q}}{q^2} \delta(\mathbf{q} \cdot \mathbf{v} - \omega). \quad (2.25)$$

We are now interested in the total energy U that is lost per unit time. The energy loss rate is given by the sum of the kinetic energy and the induced potential energy:

$$\begin{aligned} \frac{dU}{dt} &= \frac{d}{dt} \left(\frac{1}{2} m v^2 - e \phi_{\text{ind}}(\mathbf{r}, t) \right) \\ &= m\mathbf{v} \cdot \frac{d\mathbf{v}}{dt} - e \left(\frac{\partial \phi_{\text{ind}}(\mathbf{r}, t)}{\partial t} + \sum_{i=1}^3 \frac{\partial \phi_{\text{ind}}(\mathbf{r}, t)}{\partial x_i} \frac{\partial x_i}{\partial t} \right) \\ &= -e \frac{\partial \phi_{\text{ind}}(\mathbf{r}, t)}{\partial t} \end{aligned} \quad (2.26)$$

where the x_i are x , y and z .

The induced electric field is obtained by subtracting the electric field that the electron would create in vacuum from the total electric field (2.25):

$$\mathbf{E}(\mathbf{q}, \omega) = \frac{ie}{2\pi\epsilon_0} \left[\frac{1}{\epsilon(\mathbf{q}, \omega)} - 1 \right] \frac{\mathbf{q}}{q^2} \delta(\mathbf{q} \cdot \mathbf{v} - \omega). \quad (2.27)$$

The induced potential can be found in Fourier space from equation (2.24) and transformed back to real space to give the rate of energy loss

$$\frac{dU}{dt} = \frac{ie^2}{(2\pi)^3\epsilon_0} \int d\omega \int d\mathbf{q} \frac{\omega}{q^2} \left[\frac{1}{\epsilon(\mathbf{q}, \omega)} - 1 \right] e^{i(\mathbf{q}\cdot\mathbf{v}-\omega)t} \delta(\mathbf{q} \cdot \mathbf{v} - \omega). \quad (2.28)$$

We assume that the medium is homogeneous and isotropic: $\epsilon(\mathbf{q}, \omega) = \epsilon(q, \omega)$. The angular part of the \mathbf{q} integral can then be performed: the azimuthal integral yields 2π and the polar integral can be solved using the delta function. The result is

$$\frac{dU}{dt} = \frac{ie^2}{(2\pi)^2\epsilon_0} \int_{-qv}^{qv} d\omega \int_0^\infty dq \frac{\omega}{qv} \left[\frac{1}{\epsilon(q, \omega)} - 1 \right]. \quad (2.29)$$

With the property that $\epsilon(\mathbf{q}, \omega) = \epsilon^*(-\mathbf{q}, -\omega)$, the real part of the integrand is anti-symmetric in ω and therefore integrates to zero. The imaginary part is symmetric, so

$$\frac{dU}{dt} = -\frac{e^2}{2\pi^2\epsilon_0} \int_0^\infty dq \int_0^{qv} d\omega \frac{\omega}{qv} \text{Im} \left[-\frac{1}{\epsilon(q, \omega)} \right]. \quad (2.30)$$

A more interesting quantity is the stopping power S , the rate of energy loss per unit path length:

$$\begin{aligned} S &= -\frac{dU}{ds} = -\frac{1}{v} \frac{dU}{dt} \\ &= \frac{e^2}{2\pi^2\epsilon_0 v^2} \int \frac{dq}{q} \int d\omega \omega \text{Im} \left[-\frac{1}{\epsilon(q, \omega)} \right]. \end{aligned} \quad (2.31)$$

Discrete stopping events

In the previous section, the stopping of the electron has been considered as a continuous process. In reality, this is not true: the electron undergoes a series of discrete events in which a finite amount of energy is lost. The translation between the two descriptions is fairly straightforward. The Fourier variables ω and q are associated with energy transfer, $\hbar\omega$, and momentum transfer, $\hbar q$, respectively.

In a full quantum mechanical calculation, one would find cross sections or inverse mean free paths λ^{-1} . The probability for the electron to lose energy $\hbar\omega$ and momentum $\hbar\mathbf{q}$ would be given by the differential inverse mean free path,

$$\frac{d^2\lambda^{-1}}{d(\hbar\omega) d(\hbar\mathbf{q})}.$$

The stopping power is defined as the expected amount of energy lost by the electron,

$$S = \int d(\hbar\omega) \hbar\omega \int d(\hbar\mathbf{q}) \frac{d^2\lambda^{-1}}{d(\hbar\omega) d(\hbar\mathbf{q})},$$

where the limits of integration are given by what is kinematically allowed (see section 2.1.1). By association with equation (2.31), we can read off the differential inverse mean free path to be

$$\frac{d^2\lambda^{-1}}{d\omega dq} = \frac{e^2}{2\pi^2\epsilon_0\hbar v^2} \frac{1}{q} \operatorname{Im} \left[-\frac{1}{\epsilon(q, \omega)} \right]. \quad (2.32)$$

Despite the fact that the transverse components of the electromagnetic field were neglected, it is common [26] to relate v to the electron's kinetic energy T by means of the relativistic energy-momentum relationship

$$v^2 = \frac{2T}{m} \frac{1 + T/2mc^2}{(1 + T/mc^2)^2}. \quad (2.33)$$

The mean free path between interactions then becomes

$$\lambda^{-1} = \frac{\hbar}{\pi a_0 T} \frac{(1 + T/mc^2)^2}{1 + T/2mc^2} \int d\omega \int \frac{dq}{q} \operatorname{Im} \left[-\frac{1}{\epsilon(q, \omega)} \right]. \quad (2.34)$$

a_0 is the Bohr radius. Integration is over all kinematically allowed values of q and ω . The energy loss is limited to $0 \leq \hbar\omega \leq T$; the limits on q are given by equation (2.10).

The quantity $\operatorname{Im}[-1/\epsilon(q, \omega)]$ is the only material-specific factor in this equation. It represents the probability of losing a certain amount of energy and momentum. In literature, it is also known as the energy-loss function. An accurate description of this parameter is very important in order to obtain accurate results.

2.4.2 Obtaining the dielectric function

The energy-loss function, $\operatorname{Im}[-1/\epsilon(q, \omega)]$, can be obtained in multiple ways.

One possibility is density functional theory [27, 28]. Unfortunately, this is computationally very expensive, and the unit cell is limited to only a few atoms. Furthermore, only small values of the (q, ω) parameters can be used.

By far the most popular method is to measure $\epsilon(0, \omega)$ at $q = 0$, and then fit a free-electron gas model to extrapolate it for $q > 0$. $\epsilon(0, \omega)$ can be measured with light as the refractive index n and the extinction coefficient k ,

$$\epsilon = (n + ik)^2, \quad (2.35)$$

where n and k implicitly depend on ω , the energy of the photon. n and k values have been measured for a wide range of materials [29–32]. In the infrared and




Figure 2.6: The energy-loss function for silicon. Data was combined from various sources [29, 30, 34, 35].

visible part of the wavelength spectrum, n and k values may be obtained by Fourier-transform infrared spectroscopy and ellipsometry. In the x-ray range, the atomic scattering factors may be calculated as the photons are primarily sensitive to deep inner-shell electrons. In the intermediate range, approximately between 6–30 eV, it is often difficult to obtain accurate data. Synchrotrons are the most accurate option, but they are expensive.

The energy-loss function can also be measured by electron energy loss spectroscopy, usually in transmission through thin films [33]. Transmitted electrons that have undergone inelastic interaction have both different energies and directions, which, in principle, allows exploration of both q and ω at once. It is difficult to measure large q values, because large-angle scattering will cause the incident particles to undergo multiple scattering interactions. The energy-loss function can also be measured in reflection [31]. One must, however, be very careful to eliminate contributions from multiple scattering.

It is now instructive to discuss the main features of the energy-loss function. Figure 2.6 shows the energy-loss function of silicon. Two sharp ridges appear near 100 eV and 1.8 keV. These represent the binding energies of the inner K and L electron shells of Si atoms. Photons with less energy than an electron's binding energy cannot excite this electron. When they have more energy, this absorption channel opens, resulting in a sharp edge in the energy-loss function.

A silicon atom has four electrons outside of the K and L shells. In the solid state, these electrons are shared between atoms, forming bonds. Their energy levels are no longer discrete, but instead they form continuous bands and be-

have similarly to a free-electron gas. The peak at 17 eV corresponds to the resonant frequency of this free-electron gas, also known as the bulk plasmon. Excitation of a bulk plasmon is by far the dominant energy loss channel for charged particles in a solid.

Below 1 eV, the energy-loss function drops to practically zero. This is due to the band gap of silicon: an electron from the bulk cannot be excited to a state in this forbidden gap. At very low energies, resonant peaks appear again. These are not electronic excitations, but vibrational modes of the lattice: longitudinal optical phonons.

The energy-loss function captures a wide variety of physics. This is one of the reasons for its popularity for describing electron-matter interaction. However, because photons have negligible momentum, it is usually only measured in the “optical” ($q \approx 0$) regime. After $\epsilon(0, \omega)$ has been obtained, a model needs to be fitted to find the dielectric function for arbitrary momentum transfers. It is clear from figure 2.6 that the dominant loss channel for silicon is the bulk plasmon. The same is true for all other materials: interactions with valence electrons (either direct excitations of electron-hole pairs or plasmons) dominate over inner-shell electrons and phonons. Since valence electrons can be well described as a free-electron gas, the dielectric function for a free-electron gas is the most popular type of fit function.

2.4.3 The Lindhard dielectric function

We will now derive the Lindhard dielectric function [36] for a free electron gas. Rigorous derivations are available in literature [37], but the goal here is to gain an intuitive understanding of the physics that is included.

If we apply an external potential, the charges in the electron gas will rearrange themselves, producing an induced potential. This acts to screen the external potential, modifying the total potential in the system. The problem is to solve the system self-consistently: the electrons in the gas respond to the total potential, including the screening effect of their surroundings.

Even for a model as simple as an electron gas, no exact expression for the dielectric function exists. The reason is that it is difficult to account for interactions between the electrons in the gas. In reality, each electron is surrounded by a reduced electronic density because it repels its neighbours by the Coulomb force and the exchange interaction. We assume that this effect is negligible, and that electrons respond only to the average field generated by all other electrons. This approximation is only valid in the limit of high electron densities. In prac-

tice, this approximation means that if we apply an external potential $V(\mathbf{q}, \omega)$ (oscillating both in time and space), the electrons in the gas will oscillate in phase with the potential. Any out-of-phase response, which depends on the position of the electrons, is assumed to average out. This is called the random-phase approximation.

We make two further assumptions. As before, we neglect the magnetic field and work in the Coulomb gauge. This makes the equations formally electrostatic. The second assumption is that both the external and induced potentials are weak, allowing the use of first-order perturbation theory.

The electric field \mathbf{E} is generated by the sum of “free” charges ρ_f and induced charges ρ_i :

$$\rho(\mathbf{r}, t) = \rho_f(\mathbf{r}, t) + \rho_i(\mathbf{r}, t), \quad (2.36)$$

$$\nabla \cdot \mathbf{E}(\mathbf{r}, t) = \rho(\mathbf{r}, t)/\epsilon_0. \quad (2.37)$$

The fields can be written in terms of scalar potentials:

$$\mathbf{E}(\mathbf{r}, t) = -\nabla\phi(\mathbf{r}, t), \quad (2.38a)$$

$$\mathbf{D}(\mathbf{r}, t) = -\epsilon_0\nabla\phi_f(\mathbf{r}, t), \quad (2.38b)$$

which are related to the charge density by

$$\rho(\mathbf{r}, t) = -\epsilon_0\nabla^2\phi(\mathbf{r}, t), \quad (2.39a)$$

$$\rho_f(\mathbf{r}, t) = -\epsilon_0\nabla^2\phi_f(\mathbf{r}, t). \quad (2.39b)$$

In Fourier space,

$$\rho(\mathbf{q}, \omega) = \epsilon_0q^2\phi(\mathbf{q}, \omega), \quad (2.40a)$$

$$\rho_f(\mathbf{q}, \omega) = \epsilon_0q^2\phi_f(\mathbf{q}, \omega). \quad (2.40b)$$

With the same linear relationship (2.21) between the electric field and the displacement field, the dielectric function can be associated with

$$\epsilon(\mathbf{q}, \omega) = 1 - \frac{1}{\epsilon_0q^2} \frac{\rho_i(\mathbf{q}, \omega)}{\phi(\mathbf{q}, \omega)}. \quad (2.41)$$

We will consider plane-wave eigenstates of the unperturbed Hamiltonian, $\psi_{\mathbf{k}}^{(0)}(\mathbf{r}) = e^{i\mathbf{k}\cdot\mathbf{r}}/\sqrt{\Omega}$, where Ω is the volume of space. With the random-phase approximation in mind, we will derive the system’s response to the following (external + induced) perturbation to the Hamiltonian:

$$V(\mathbf{r}, t) = \left(V e^{i(\mathbf{q}\cdot\mathbf{r}-\omega t)} + V^* e^{-i(\mathbf{q}\cdot\mathbf{r}-\omega t)} \right) e^{\gamma t}. \quad (2.42)$$

γ has been introduced as an infinitesimally small positive constant, ensuring that the perturbation is slowly switched on. In the end, the limit $\gamma \downarrow 0$ will be taken.

In first-order perturbation theory, the perturbed wave functions, $\psi_{\mathbf{k}}$, can be expressed as

$$\psi_{\mathbf{k}}(\mathbf{r}, t) = \sum_{\mathbf{k}'} C_{\mathbf{k},\mathbf{k}'} \psi_{\mathbf{k}}^{(0)}(\mathbf{r}) e^{-iE_{\mathbf{k}'}t/\hbar}, \quad (2.43)$$

where

$$C_{\mathbf{k},\mathbf{k}'} = -\frac{i}{\hbar} \int_{-\infty}^t dt' \langle \psi_{\mathbf{k}'}^{(0)} | V | \psi_{\mathbf{k}}^{(0)} \rangle e^{i(E_{\mathbf{k}'} - E_{\mathbf{k}})t'/\hbar}. \quad (2.44)$$

There are only two non-zero $C_{\mathbf{k},\mathbf{k}'}$: those for which $\mathbf{k}' = \mathbf{k} \pm \mathbf{q}$. Specifically, they are

$$C_{\mathbf{k},\mathbf{k}+\mathbf{q}} = \frac{V e^{i(E_{\mathbf{k}+\mathbf{q}} - E_{\mathbf{k}})t/\hbar} e^{-i\omega t} e^{\gamma t}}{E_{\mathbf{k}} - E_{\mathbf{k}+\mathbf{q}} + \hbar\omega - i\hbar\gamma}, \quad (2.45a)$$

$$C_{\mathbf{k},\mathbf{k}-\mathbf{q}} = \frac{V^* e^{i(E_{\mathbf{k}-\mathbf{q}} - E_{\mathbf{k}})t/\hbar} e^{i\omega t} e^{\gamma t}}{E_{\mathbf{k}} - E_{\mathbf{k}-\mathbf{q}} - \hbar\omega + i\hbar\gamma}. \quad (2.45b)$$

We now obtain the charge density, ρ , of the electron gas, by summing over the squares of all perturbed wave functions $\psi_{\mathbf{k}}$. Subtracting the unperturbed $|\psi_{\mathbf{k}}^{(0)}|^2$ yields the induced charge. We denote the initial occupation of each state by $f(\mathbf{k})$, which will be two times the Fermi distribution (the factor two accounting for the spin degeneracy). Keeping only terms linear in V , one finds for the induced charge density

$$\begin{aligned} \rho_i(\mathbf{r}, t) &= e \sum_{\mathbf{k}} f(\mathbf{k}) \left(|\psi_{\mathbf{k}}(\mathbf{r}, t)|^2 - |\psi_{\mathbf{k}}^{(0)}(\mathbf{r})|^2 \right) \\ &= e \sum_{\mathbf{k}} f(\mathbf{k}) \left(\frac{V e^{i(\mathbf{q}\cdot\mathbf{r} - \omega t)} e^{\gamma t}}{E_{\mathbf{k}} - E_{\mathbf{k}+\mathbf{q}} + \hbar\omega + i\hbar\gamma} + \frac{V^* e^{-i(\mathbf{q}\cdot\mathbf{r} - \omega t)} e^{\gamma t}}{E_{\mathbf{k}} - E_{\mathbf{k}-\mathbf{q}} - \hbar\omega + i\hbar\gamma} + \text{c.c.} \right), \end{aligned} \quad (2.46)$$

where c.c. stands for the complex conjugate of the first two terms. In the second term (and its complex conjugate), the dummy summation variable \mathbf{k} can be replaced by $\mathbf{k} + \mathbf{q}$. Collecting the terms,

$$\rho_i(\mathbf{r}, t) = e \sum_{\mathbf{k}} \left(\frac{f(\mathbf{k}) - f(\mathbf{k} + \mathbf{q})}{E_{\mathbf{k}} - E_{\mathbf{k}+\mathbf{q}} + \hbar\omega + i\hbar\gamma} V e^{i(\mathbf{q}\cdot\mathbf{r} - \omega t)} e^{\gamma t} + \text{c.c.} \right). \quad (2.47)$$

The Fourier components $\rho(\mathbf{q}, \omega)$ are easily found. With $V(\mathbf{q}, \omega) = e\phi(\mathbf{q}, \omega)$,

$$\rho_i(\mathbf{q}, \omega) = \frac{e^2}{\epsilon_0 q^2} \sum_{\mathbf{k}} \frac{f(\mathbf{k}) - f(\mathbf{k} + \mathbf{q})}{E_{\mathbf{k}} - E_{\mathbf{k}+\mathbf{q}} + \hbar\omega + i\hbar\gamma} \phi(\mathbf{q}, \omega). \quad (2.48)$$

Therefore, referring to equation (2.41),

$$\epsilon(\mathbf{q}, \omega) = 1 - \frac{e^2}{\epsilon_0 q^2} \sum_{\mathbf{k}} \frac{f(\mathbf{k}) - f(\mathbf{k} + \mathbf{q})}{E_{\mathbf{k}} - E_{\mathbf{k} + \mathbf{q}} + \hbar\omega + i\hbar\gamma}. \quad (2.49)$$

This is the famous Lindhard dielectric function [36] as it is often found in textbooks. The wavevector \mathbf{k} runs over all space. It is convenient to substitute $\mathbf{k} \rightarrow -\mathbf{k} - \mathbf{q}$ in the $f(\mathbf{k} + \mathbf{q})$ term. Then, using that $f(-\mathbf{k}) = f(\mathbf{k})$ and $E_{-\mathbf{k}} = E_{\mathbf{k}}$,

$$\epsilon(\mathbf{q}, \omega) = 1 - \frac{e^2}{\epsilon_0 q^2} \sum_{\mathbf{k}} \left(\frac{f(\mathbf{k})}{E_{\mathbf{k}} - E_{\mathbf{k} + \mathbf{q}} + \hbar\omega + i\hbar\gamma} + \frac{f(\mathbf{k})}{E_{\mathbf{k}} - E_{\mathbf{k} + \mathbf{q}} - \hbar\omega - i\hbar\gamma} \right) \quad (2.50)$$

$$= 1 - \frac{2e^2}{\epsilon_0 q^2} \sum_{\mathbf{k}} f(\mathbf{k}) \frac{E_{\mathbf{k}} - E_{\mathbf{k} + \mathbf{q}}}{[E_{\mathbf{k}} - E_{\mathbf{k} + \mathbf{q}}]^2 - [\hbar\omega + i\hbar\gamma]^2}. \quad (2.51)$$

This form is convenient for physical interpretation, since each term represents the contribution of a particular electronic state.

Explicit form at zero temperature

We will now find an explicit form of the Lindhard dielectric function at zero temperature. In this case, the Fermi distribution $f(\mathbf{k})$ is a step function, equal to 1 for $|\mathbf{k}| < |\mathbf{k}_F|$ and zero otherwise. k_F is the Fermi wave vector.

We start from equation (2.50) and replace the sum over \mathbf{k} by an integral.

$$\epsilon(\mathbf{q}, \omega) = 1 - \frac{e^2}{\epsilon_0 q^2} \int \frac{d\mathbf{k}}{(2\pi)^3} \left(\frac{f(\mathbf{k})}{E_{\mathbf{k}} - E_{\mathbf{k} + \mathbf{q}} + \hbar\omega + i\hbar\gamma} + \frac{f(\mathbf{k})}{E_{\mathbf{k}} - E_{\mathbf{k} + \mathbf{q}} - \hbar\omega - i\hbar\gamma} \right). \quad (2.52)$$

We now use the identity

$$\frac{1}{x \pm i\gamma} = \mathcal{P} \frac{1}{x} \mp i\pi\delta(x), \quad (2.53)$$

where \mathcal{P} stands for the principal part of the integral and $\delta(x)$ is the Dirac delta function. It allows us to split the dielectric function into a real and imaginary part:

$$\epsilon(\mathbf{q}, \omega) = \epsilon_1(\mathbf{q}, \omega) + i\epsilon_2(\mathbf{q}, \omega), \quad (2.54a)$$

where

$$\epsilon_1(\mathbf{q}, \omega) = 1 - \frac{e^2}{\epsilon_0 q^2} \mathcal{P} \int \frac{d\mathbf{k}}{(2\pi)^3} \left(\frac{f(\mathbf{k})}{E_{\mathbf{k}} - E_{\mathbf{k} + \mathbf{q}} + \hbar\omega} + \frac{f(\mathbf{k})}{E_{\mathbf{k}} - E_{\mathbf{k} + \mathbf{q}} - \hbar\omega} \right), \quad (2.54b)$$

$$\epsilon_2(\mathbf{q}, \omega) = \frac{\pi e^2}{\epsilon_0 q^2} \int \frac{d\mathbf{k}}{(2\pi)^3} f(\mathbf{k}) \left[\delta(\hbar\omega + E_{\mathbf{k}} - E_{\mathbf{k} + \mathbf{q}}) - \delta(\hbar\omega - E_{\mathbf{k}} + E_{\mathbf{k} + \mathbf{q}}) \right]. \quad (2.54c)$$

To find ϵ_1 , we perform the integral in spherical coordinates. Abbreviating the cosine of the polar integration angle as $c = \cos\theta$, we find that each term contributes

$$\begin{aligned} \int \frac{d\mathbf{k}}{(2\pi)^3} \frac{1}{E_{\mathbf{k}} - E_{\mathbf{k}+\mathbf{q}} \pm \hbar\omega} &= \frac{1}{(2\pi)^2} \int_0^{k_F} dk k^2 \int_{-1}^1 dc \frac{2m}{\pm 2m\hbar\omega - \hbar^2 q^2 - 2\hbar^2 k q c} \\ &= \frac{m}{(2\pi\hbar)^2} \int_0^{k_F} dk \frac{k}{q} \ln \left| \frac{-q^2 + 2kq \pm 2m\omega/\hbar}{-q^2 - 2kq \pm 2m\omega/\hbar} \right| \\ &= \frac{mk_F^2}{8\pi^2\hbar^2 q} \left[\left\{ 1 - (z \mp u)^2 \right\} \ln \left| \frac{z \mp u - 1}{z \mp u + 1} \right| - 2(z \mp u) \right]. \end{aligned}$$

The substitutions $z = q/(2k_F)$ and $u = \omega/(qv_F)$ were made in the last step, where $v_F = \hbar k_F/m$ is the Fermi velocity.

ϵ_2 can also be found by integration in spherical coordinates. Because of the delta functions, the integrand is only finite for $\hbar\omega = E_{\mathbf{k}+\mathbf{q}} - E_{\mathbf{k}}$. Using the delta function for integrating out the polar angle, each term contributes

$$\begin{aligned} \int \frac{d\mathbf{k}}{(2\pi)^3} \delta(\hbar\omega \pm E_{\mathbf{k}} \mp E_{\mathbf{k}+\mathbf{q}}) &= \frac{1}{(2\pi)^2} \int_0^{k_F} dk k^2 \int_{-1}^1 dc \delta\left(\hbar\omega \mp \frac{\hbar^2 q^2}{2m} \mp \frac{\hbar^2 k q c}{m}\right) \\ &= \frac{m}{(2\pi\hbar)^2 q} \int_0^{k_F} dk k \Theta\left(-k < \frac{m\omega}{\hbar q} \mp \frac{q}{2} < k\right) \\ &= \frac{mk_F^2}{8\pi^2\hbar^2 q} \left\{ 1 - (u \mp z)^2 \right\} \Theta(-1 < u \mp z < 1). \end{aligned}$$

The function Θ is the “logical step function”, a natural extension of the Heaviside step function that is equal to one when its argument is true and zero otherwise.

Putting everything together, we find that the Lindhard dielectric function at zero temperature is given by the following beauty:

$$\epsilon(\mathbf{q}, \omega) = \epsilon_1(\mathbf{q}, \omega) + i\epsilon_2(\mathbf{q}, \omega), \quad (2.55a)$$

$$\begin{aligned} \epsilon_1(\mathbf{q}, \omega) &= 1 + \frac{2}{\pi a_0 q z} \left[\frac{1}{8z} \left\{ 1 - (z - u)^2 \right\} \ln \left| \frac{z - u + 1}{z - u - 1} \right| \right. \\ &\quad \left. + \frac{1}{8z} \left\{ 1 - (z + u)^2 \right\} \ln \left| \frac{z + u + 1}{z + u - 1} \right| + \frac{1}{2} \right], \end{aligned} \quad (2.55b)$$

$$\epsilon_2(\mathbf{q}, \omega) = \frac{1}{4a_0 q z^2} \begin{cases} 4uz & \text{for } z + u < 1, \\ 1 - (z - u)^2 & \text{for } |z - u| < 1 < z + u, \\ 0 & \text{otherwise,} \end{cases} \quad (2.55c)$$

where $z = q/(2k_F)$, $u = \omega/(qv_F)$, and $a_0 = 4\pi\epsilon_0/m e^2$ is the Bohr radius.

Physical content of the Lindhard function

In equation (2.51), each term in the sum over \mathbf{k} represents the contribution of that electronic state to the dielectric function. There is a singularity when $E_{\mathbf{k}} - E_{\mathbf{k}+\mathbf{q}} = \hbar\omega$. This pole gave rise to the imaginary component of the dielectric function, equation (2.54c), and represents direct excitation of an electron from \mathbf{k} to $\mathbf{k} + \mathbf{q}$. This process is also called “single excitation”.

There is a second effect, when $\hbar\omega$ far exceeds the energy differences between electronic states. Using that

$$\sum_{\mathbf{k}} f(\mathbf{k}) = n/2 \quad (2.56a)$$

and

$$\sum_{\mathbf{k}} f(\mathbf{k}) \left(\frac{\hbar k}{m} \right)^2 = \frac{3}{5} v_F^2 \frac{n}{2}, \quad (2.56b)$$

where n is the number density of electrons, a Taylor expansion of equation (2.51) around $q = 0$ gives

$$\epsilon(q, \omega) \approx 1 - \frac{\hbar^2 \omega_p^2}{(\hbar\omega + i\hbar\gamma)^2} - \frac{3}{5} \frac{\hbar^2 \omega_p^2}{(\hbar\omega + i\hbar\gamma)^4} v_F^2 q^2 - \dots, \quad (2.57)$$

where ω_p is defined by

$$\omega_p^2 = \frac{ne^2}{m\epsilon_0}. \quad (2.58)$$

The dielectric function is close to zero when ω is close to ω_p . More precisely, it happens when

$$\omega(q) = \omega_p \left(1 + \frac{3}{10} \frac{v_F^2 q^2}{\omega_p^2} + \dots \right). \quad (2.59)$$

If the dielectric function is close to zero, a small external field gives rise to a large induced field. ω_p is equal to the natural eigenfrequency of an electron gas. An excitation on the dispersion line of equation (2.59) represents a collective resonant oscillation of all electrons in the gas. ω_p is called the plasma frequency, and a single quantum of the plasma resonance is called a plasmon. The energy-loss function contains the plasmon peak as a sharp delta function. It can be shown that, at $q = 0$,

$$\lim_{q \rightarrow 0} \text{Im} \frac{-1}{\epsilon(q, \omega)} = \frac{\pi}{2} \omega_p \delta(\omega - \omega_p). \quad (2.60)$$




Figure 2.7: Regions in energy-momentum space with different types of excitation contained in the Lindhard dielectric function. The shaded area is the region of single excitation, the solid curve is the plasma resonance. The dashed curve represents the boundary between the two regions in equation (2.55c) where the imaginary part is nonzero. This graph was made for $\hbar\omega_p = 15$ eV.

Figure 2.7 illustrates the regions of energy-momentum space where these types of excitations are possible. For the convenience of the reader, this figure uses the auxiliary variable $x = \hbar\omega/E_F = 4uz$ instead of u . Single excitation is possible in the shaded region, the plasma resonance is given by the black curve.

Figure 2.7 shows that the plasmon curve intersects the single-excitation region at large momentum transfers. This gives rise to so-called Landau damping, in which a plasmon decays by transferring all energy to a single electron. However, the Lindhard dielectric function features no decay mechanism for plasmons at lower momentum transfers. The plasmon lifetime at small q is therefore infinite, and the Lindhard energy-loss function is infinitely sharp around the plasmon peak. In real materials, plasmons decay via a variety of processes, such as scattering of the excited electrons on impurities, or the production of multiple electron-hole pairs.

It is tempting to introduce plasmon damping at small q values by allowing finite values for the infinitesimal constant γ . This prevents ϵ from equalling zero in equation (2.57) when $\omega \approx \omega_p$. In fact, Lindhard originally suggested this approach [36], but it was later shown [38] that this is not correct because it fails to conserve the local electron number. A more correct approach was proposed by

Mermin [39]. His dielectric function takes the form

$$\epsilon_M(\mathbf{q}, \omega) = 1 + \frac{(1 + i\gamma/\omega)(\epsilon_L(\mathbf{q}, \omega + i\gamma) - 1)}{1 + (i\gamma/\omega)(\epsilon_L(\mathbf{q}, \omega + i\gamma) - 1)/(\epsilon_L(\mathbf{q}, 0) - 1)}. \quad (2.61)$$

Here, $\epsilon_L(\mathbf{q}, \omega)$ is the Lindhard dielectric function. γ is interpreted as the inverse lifetime of the plasmon. The value of γ is not predicted by the Mermin model, so this model has two free parameters.

We mention one further model, which is called the plasmon pole approximation. The plasmon pole approximation describes a system in which the dynamics are entirely due to the collective plasma oscillation. It is not a specific limit to the Lindhard dielectric function. This model dielectric function reads

$$\epsilon_{PP}(q, \omega) = 1 + \frac{\omega_p^2}{\omega_p^2(q) - \omega_p^2 - \omega^2 - i\omega\gamma}. \quad (2.62)$$

Here, $\omega_p(q)$ is the dispersion relation for the plasma excitation, usually given by equation (2.59). A damping factor γ is explicitly included in the plasmon pole approximation. At $q = 0$, this dielectric function is equivalent to the famous Drude dielectric function. The plasmon pole approximation is sometimes called the extended Drude model in literature. In the limit of zero damping, its energy-loss function becomes a delta function,

$$\lim_{\gamma \rightarrow 0} \text{Im} \frac{-1}{\epsilon_{PP}(q, \omega)} = \frac{\pi}{2} \frac{\omega_p^2}{\omega_p(q)} \delta(\omega - \omega_p(q)). \quad (2.63)$$

2.4.4 Optical data models

As mentioned before, it is popular to measure the dielectric function in the optical ($q = 0$) regime, and then extend it to non-zero momenta by fitting a model dielectric function. It appears that the idea to use measured optical data was first used by Howie and Stern [40], while Ritchie and Howie [41] were the first to fit a dielectric function model at $q = 0$. They use a linear combination of multiple plasmon-pole oscillators of the form of equation (2.62).

Penn [42] suggested a model where the extension is preformed using the Lindhard dielectric function. While the model of Ritchie and Howie [41] fits a sum over an arbitrary number of Drude oscillators, the Penn model can be understood as fitting an infinite number of Lindhard functions:

$$\text{Im} \frac{-1}{\epsilon(q, \omega)} = \int_0^\infty d\omega_p g(\omega_p) \text{Im} \frac{-1}{\epsilon_L(q, \omega, \omega_p)}. \quad (2.64)$$

Here, $\epsilon_L(q, \omega, \omega_p)$ is the Lindhard dielectric function with an electron density such that the plasma energy is ω_p (see equation (2.58)). The function $g(\omega_p)$ weights the amplitude for each Lindhard oscillator. Using equation (2.60) and the requirement that $\text{Im} \frac{-1}{\epsilon(q, \omega)}$ equals the measured optical data at $q = 0$, it is easy to see that

$$g(\omega) = \frac{2}{\pi \omega} \text{Im} \frac{-1}{\epsilon(\omega)}, \quad (2.65)$$

where $\epsilon(\omega)$ is the measured optical data.

Penn also suggests a similar algorithm which swaps the Lindhard dielectric function for the zero-damping plasmon pole approximation of equation (2.63). This provides similar results at high electron energies of $\gtrsim 200$ eV. Penn uses the dispersion relation of equation (2.59). A similar approach is taken by Ashley [43, 44] and independently by Ding and Shimizu [45], except they use the approximate dispersion relation

$$\hbar\omega_p(q) = \hbar\omega_p + \frac{\hbar^2 q^2}{2m}. \quad (2.66)$$

This significantly reduces the computational difficulty. In Ashley's model, the integration over q to find the mean free path and stopping power can be done analytically. Another advantage comes in simulations of electron scattering, where the joint probability distribution of both energy and momentum loss is required. This normally means that a 2-dimensional table needs to be stored. If Ashley's model is used, only a single variable needs to be stored; the other random variable can be sampled analytically. One major drawback of Ashley's model is the unphysical property that the minimum kinetic energy required to excite an electron is shifted to twice that electron's binding energy. This is a problem for inner shells, which have binding energies that are often similar to the energy of the incident electron. Kieft and Bosch [10], who built an electron scattering simulator based on Ashley's model, needed to add semi-empirical corrections to account for this fact.

Another extension mechanism is to fit a finite number of Mermin oscillators [46]. Because the extended Drude model and the Mermin model are identical at $q = 0$, the fitting parameters from the method of Ritchie and Howie [41] can be used directly. As the extension to $q > 0$ is provided by the more sophisticated Mermin function, the results should be more accurate. However, it is not always possible to get a good fit at $q = 0$, especially around sharp features in the energy-loss function such as the band gap or the sharp ridges of inner shells. Da *et al.* [47] solve this by using a very large number of Mermin oscillators and

allowing negative amplitudes. Abril *et al.* [48] use Mermin oscillators for inner shells, with a sharp cutoff at the shell's binding energy. Later, the same group [49, 50] proposed the so-called MELF-GOS model, which uses Mermin functions (MELF) for valence electrons and generalized oscillator strengths (GOS) [51] for inner shells. Essentially, the GOS represents the ionization cross section of inner electron shells. Analytical expressions [49] for the GOS of hydrogen are typically used, even for larger elements.

The last method we mention here is an interesting one proposed by Bourke and Chantler [52]. They fit an infinite number of Mermin oscillators with a q -dependent damping parameter $\gamma(\omega, q)$. They set $\gamma(\omega, q = 0) = 0$, so the Penn algorithm can be used to fit the optical energy-loss function. Then they use an iterative scheme based on the measured optical data to fit the q -dependence of the damping coefficient. Compared to the schemes mentioned above, which fit a finite number of Mermin oscillators, this scheme has the advantage of having no free parameters and being able to capture all details of the optical ELF.

We conclude this section with the mean free path and stopping power of silicon for the various models discussed here. These can be seen in figure 2.8. The Ritchie and Mermin models were computed using a single oscillator with $\omega_p = 16.87$ eV and $\gamma = 4.245$ eV. These parameters were taken from Abril *et al.* [48]. The discrepancy in stopping powers between the Ashley & Penn and Ritchie & Mermin models at high energy is due to the fact that inner shells are not present in the energy-loss function of the latter two. Inner-shell ionization events are not very likely, hence the good agreement in mean free path at high energies, but a lot of energy is lost if these events do take place. Inner-shell ionization should be taken into account in a complete model. The models diverge at < 100 eV energies.

2.4.5 Conclusion

The mean free path for inelastic scattering on the electrons in the solid is given by equation (2.34). The difficulty lies in obtaining data for the energy-loss function, $\text{Im}[-1/\epsilon(q, \omega)]$. This can be done directly by density functional theory, but it is more common to obtain $\epsilon(0, \omega)$ experimentally and fit one of the models discussed in section 2.4.4.

The model of Ashley [43] is the simplest, but produces substantially longer mean free paths at low energies than the other (more accurate) models (see figure 2.8a). It is more appropriate to use the model of Penn [42], to which Ashley's model can be seen as an approximation. The Penn model has the disadvantage




Figure 2.8: The (a) mean free path and (b) stopping power for silicon, calculated using the various models mentioned in this section. Measured data for the stopping power were taken from [53].

of being computationally more difficult, but it is very popular in modern literature [17, 54–57]. We choose to use the Penn model in this thesis.

Theoretically, one would expect it is better to use Mermin oscillators instead. After all, the Mermin dielectric function includes damping, which Lindhard’s dielectric function (on which the Penn model is based) does not. Mermin oscillators have the disadvantage that an unknown number of oscillators need to be fitted. The simulation result therefore becomes operator dependent. Mermin oscillators generally yield shorter mean free paths and higher stopping powers at low energies than the Penn model. Empirically, it seems that Mermin oscillators may overshoot the true values somewhat. This will become apparent in section 3.6, when we investigate the impact of shortening the inelastic mean free path at low energies. Ridzel *et al.* [55] have proposed interpolating between the Penn and Mermin models at low energies, tuning them to find the correct electron yields. We use different elastic scattering models at low energies in this thesis than Ridzel *et al.* did. Our elastic models generally cause electrons to scatter more, hence our “optimal” model is much closer to the Penn model than to Mermin oscillators. We therefore choose to use only the Penn model in this thesis.

The dielectric function model gives us the mean free path and the probability distributions for energy and momentum transfer when an electron scatters. It does not tell us where that energy and momentum goes. Most of the time, a secondary electron is created in such a scattering event. The dielectric func-

tion model does not tell us what the secondary electron's energy and direction will be. That requires a separate model, which will be the topic of the following section.

2.5 Secondary electron generation

Optical data models can accurately describe the mean free path between inelastic events. When an event takes place, they also give the probability distributions of energy and momentum transfer. These quantities determine the trajectory of the primary electron after the event. Optical data models do not describe what happens to the material: whether a secondary electron is excited or not, and if so, which energy and direction it gets. Additional models are needed for that, which we will describe below.

If the primary electron loses less energy than the band gap, we assume that this is due to the excitation of a longitudinal optical phonon (see also figure 2.6). The angular distribution of the primary electron is very strongly forward peaked [58]. We approximate this situation as the electron losing energy without any deflection.

In most inelastic events, a much larger amount of energy is transferred. This usually leads to the excitation of a secondary electron. We will distinguish three possibilities: direct excitation of a valence electron, excitation via plasmon decay, or inner-shell ionization. Before getting into the details of these processes, let us first discuss how the appropriate type of event is selected.

Binding energies of electrons, as well as their ionization cross sections, have been calculated for isolated, neutral cold atoms. They are publicly available in the LLNL Evaluated Electron Data Library (EEDL) [59], which was released as part of the ENDF/B library [60]. Even at finite temperature and in a solid state, such tables are quite accurate for inner shells. Based on the energy loss in the event, we probe the EEDL for the relative ionization cross sections of all electron shells. The appropriate shell can be sampled based on these cross sections. We distinguish between inner and outer shells by means of an energy threshold at 50 eV. If an inner shell is chosen, this event is performed. If an outer shell was chosen, we choose either single excitation or plasmon creation based on the energy and momentum loss. If single excitation is kinematically possible, then it is much more likely than plasmon excitation (*cf.* figure 2.7 and surrounding discussion) and single excitation is always chosen. If not enough momentum was transferred for single excitation to be kinematically possible, we assume plasmon excitation.

2.5.1 Single excitation

In a single excitation event, the primary electron transfers energy $\hbar\omega$ and momentum $\hbar\mathbf{q}$ directly to a valence electron. The primary electron's direction after the event is given by equation (2.9). What remains to be done is to choose the valence electron's energy and direction after the event.

Before the event, the valence electron has momentum $\hbar\mathbf{k}_i$ and energy $E_i = \hbar^2 k_i^2 / 2m$. It has energy $E_f = E_i + \hbar\omega$ and momentum $\hbar\mathbf{k}_f = \hbar\mathbf{k}_i + \hbar\mathbf{q}$ afterwards. Without loss of generality, we choose \mathbf{q} to be oriented along the z axis.

We have

$$\frac{\hbar^2 k_i^2}{2m} = \frac{\hbar^2}{2m} (k_x^2 + k_y^2 + k_z^2) = E_i, \quad (2.67a)$$

$$\frac{\hbar^2 k_f^2}{2m} = \frac{\hbar^2}{2m} (k_x^2 + k_y^2 + (k_z + q)^2) = E_i + \hbar\omega. \quad (2.67b)$$

This leads to

$$k_z = \frac{m\omega}{\hbar q} - \frac{q}{2}. \quad (2.68)$$

k_z is therefore determined by the optical data model, which already gives us ω and q .

k_x and k_y are constrained by the fact that $E_i < V$ and $E_f > V + G$, where V is the width of the valence band and G is the band gap (see figure 2.4). If the material is a metal, the Fermi energy takes the place of V and the band gap is zero. These constraints limit k_x and k_y to the ring

$$\frac{2m(V + G)}{\hbar^2} - (k_z + q)^2 < k_x^2 + k_y^2 < \frac{2mV}{\hbar^2} - k_z^2. \quad (2.69)$$

In simple solids, the density of states is uniformly distributed in k . We therefore choose to sample k_x and k_y uniformly from this ring. The direction and energy of the secondary electron after the event are then fully determined.

2.5.2 Plasmon excitation

After exciting a plasmon, the primary electron's direction is again given by equation (2.9). The plasmon itself may decay to a photon, or to one or multiple electron-hole pairs. A common assumption [17, 61–64] is that decay to a single electron-hole pair is the dominant mechanism, and that the secondary electron is emitted in an isotropically distributed direction. It is then assumed that the probability of choosing an electron with certain energy E_i is proportional to the

density of states before the event, multiplied by the density of states after the event. In a simple parabolic band,

$$p(E_i) \propto \sqrt{E_i(E_i + \hbar\omega)}. \quad (2.70)$$

2.5.3 Inner-shell excitation

Most optical data models aim to describe properties of valence electrons, because incoming electrons are far more likely to excite plasmons than to ionize inner shells. However, when an inner shell event happens, it involves a much larger amount of energy transfer. It is therefore important to accurately model these events, too.

In the cases of single excitation and plasmon creation, the scattering angle for the incoming electron was determined by the optical data model. For inner-shell excitation, we choose to ignore the momentum transfer given by the optical data model, but keep the probability distribution for energy transfer. We may ignore the momentum transfer because the optical data models are for free electron gases: they describe the valence electrons very well, but not the inner shells. A dedicated inner-shell model is therefore more appropriate.

First, we must choose the appropriate binding energy for the secondary electron. Electron-ionization cross sections for inner shells are publicly available in the EEDL [59, 60]. For each primary electron energy and each subshell, this library tabulates the ionization cross section and the probability distribution of energy transfer. In our case, the optical data model has already provided the energy transfer. It is a simple matter to transform the EEDL data to give the probabilities per subshell if the primary electron's energy and the energy transfer are known.

An alternative approach was taken by Kieft and Bosch [10]. Instead of electron ionization, they used photoionization cross sections, interpreting the photon energy as the energy lost by the electron. This approach has the advantage of simplicity: it does not require re-interpretation of the data and the simulator needs a smaller table (2 dimensions instead of 3 dimensions): energy loss and probability instead of primary electron energy, energy loss and probability. The trends are also correct: the innermost accessible shell, given the energy loss, is the most likely to be ionized. The relative probabilities of the subshells are approximately correct. It is worth noting that Verduin [12] introduced an error at this point: he took a similar approach to Kieft and Bosch [10] but used electron ionization cross sections instead of photoionization. This gives precisely the opposite trend, because for any given energy of the primary electron, the

outermost (valence) electrons are most likely to be ionized. The corresponding energy loss is then also small, however. If it is already known that the electron will lose a certain amount of energy, the innermost accessible shell should dominate.

With the binding energy of the inner shell selected, what remains to be done is to choose the directions of the primary and secondary electron after the event. We use a model of Ivanchenko in the Geant4 simulation toolkit [65]. This model was copied by Kieft and Bosch [10] and also used by Verduin [12]. It can be motivated starting from an elastic¹ 2-electron collision, where the target electron is initially at rest and the projectile comes in with energy T . After the collision, the projectile has energy $T - \Delta T$ and the target has ΔT . Rewriting equation (2.9) and using equation (2.7) to relate the transferred momentum, q , to the energy of the target after the collision:

$$\cos\theta_1 = \sqrt{\frac{T - \Delta T}{T}} \sqrt{\frac{1 + T/2mc^2}{1 + (T - \Delta T)/2mc^2}}. \quad (2.71)$$

θ_1 is the deflection angle of the primary electron. Conservation of momentum gives the scattering angle of the target electron:

$$\cos\theta_2 = \sqrt{\frac{\Delta T}{T}} \sqrt{\frac{1 + T/2mc^2}{1 + \Delta T/2mc^2}}. \quad (2.72)$$

θ_2 is measured with respect to the incoming projectile's direction before the event.

These scattering angles were derived for a stationary and unbound target electron, which is not true. If the target electron has binding energy B , then on average (by the Virial theorem) it is in a potential of magnitude $2B$ and has kinetic energy B . At the instant of the collision, the kinetic energy of the incoming electron is therefore $2B$ higher than it would be in vacuum. Meanwhile, the target electron's kinetic energy is B , but the scattering angles were derived for an electron at rest. Instead of transforming into the target's frame, Ivanchenko modifies the energy transfer from ΔT to $\Delta T + B$. The energy of the secondary electron after the collision is then correct. The result of Ivanchenko's model is to use equations (2.71) and (2.72), with the substitutions

$$T \rightarrow T - F + 2B, \quad (2.73a)$$

$$\Delta T \rightarrow \Delta T + B. \quad (2.73b)$$

¹The term "elastic" here means that the *total* amount of kinetic energy is conserved, not the kinetic energy of the projectile.

F is the Fermi energy. This is most likely subtracted to account for the fact that B is referenced with respect to the vacuum level while T is referenced with respect to the bottom of the band. If that is the case, however, the appropriate quantity to subtract would be the inner potential U and not the Fermi energy. It does not matter much for the final results, because F and U are on the order of a few eV, while B is on the order of a hundred eV or more.

2.6 Mott scattering

Mott scattering is a type of (quasi-)elastic scattering. The general picture is that an incoming electron is deflected by the electric field around an atom in the solid. The atom itself experiences a recoil force and a small amount of energy is transferred from the electron to the atom. However, since the electron is many orders of magnitude lighter than the atom, the energy transfer is extremely small.

The most simplistic quantitative method to describe this kind of scattering is (screened) Rutherford scattering. Rutherford scattering describes scattering due to the Coulomb interaction between two charged particles (the electron and the atomic nucleus). This ignores the electron cloud around the atom, which can be simplistically modelled by an exponentially decaying potential:

$$V(r) = \frac{Z e^2}{4\pi\epsilon_0} \frac{e^{-\alpha r}}{r}. \quad (2.74)$$

Z is the atomic number of the atom, e is the elementary charge, and α is the screening parameter representing the size of the electron cloud. This potential is called the screened Coulomb potential or the Yukawa potential in literature. The corresponding scattering cross section can be evaluated in the Born approximation, which is a popular topic for elementary quantum mechanics textbooks [66]. Empirical expressions for the screening parameter have been developed [67], and this model has been quite popular for high electron energies [10, 15, 68], especially in situations where computer power is critical.

Mott scattering [69] is conceptually similar to screened Rutherford scattering. The electrons are described by the relativistic Dirac equation instead of the Schrödinger equation. Use of the Dirac equation is also important at non-relativistic velocities because it can take spin-orbit coupling into account. The method of partial wave analysis, which is commonly done for the Schrödinger equation but can also be done for the Dirac equation, is suitable to numerically

solve scattering problems in a central potential. It is a fairly straightforward matter [70–72] to numerically obtain the differential scattering cross sections. The important step lies in the choice of potential $V(r)$.

We use the program ELSEPA by Salvat *et al.* [72]. This program can compute scattering cross sections, and has sophisticated models for the interaction potential on board. The interaction potential is a sum of several contributions:

$$V(r) = V_{\text{st}}(r) + V_{\text{ex}}(r) + V_{\text{cp}}(r) - iW_{\text{abs}}(r). \quad (2.75)$$

Each term represents the following.

- V_{st} is the electrostatic interaction potential. It is created by a sum of the nuclear charge distribution and the charge distribution of the electron cloud. The choice of nuclear charge distribution is not relevant as long as the de Broglie wavelength of the electrons is larger than the nuclear radius. We use the default Fermi distribution here, but a point charge would be equally accurate. There are several options for the density of the electron cloud. The most sophisticated one on board is also the default: numerically pre-calculated Dirac-Fock electron densities for neutral, isolated atoms. The alternatives are analytical approximations.
- V_{ex} is the exchange potential. The exchange effect accounts for the fact that the incident electron is indistinguishable from the electrons around the atom. Under the so-called static exchange approximation, this effect can be treated as an effective potential observed by the projectile. ELSEPA includes three potentials. In our experience, the exact choice of exchange potential does not matter much for the final results. We use the Furness-McCarthy potential here, which is the default option.
- V_{cp} is the correlation-polarization potential. The incident projectile polarizes the charge cloud around the target atom. This leads to a net induced dipole moment in the electric field around the atom. This effect gets stronger for slow projectiles. ELSEPA's default behaviour is to ignore this effect, but we choose the elaborate local density approximation.
- W_{abs} is an absorption term. It accounts for the fact that inelastic scattering may occur with the electrons of the atom. The effect gets stronger for slower projectiles. The main consequence is that the differential cross section is reduced for scattering angles larger than about 30° . ELSEPA has two options: on or off. This effect will be turned on in this thesis.

The potential described here applies only to isolated atoms. To describe scattering in solids, the electron density needs to be modified. If the solid is a regular crystal, the density of atoms found by the Dirac-Fock method should be recomputed using appropriate boundary conditions. Instead, ELSEPA employs the muffin-tin model. It is assumed that the interaction with an atom in the solid is modified only by its nearest neighbours. The electrons belonging to an atom are confined to within the muffin-tin radius, R_{mt} , which is half the average interatomic distance. The electron density around an atom, ρ_e , is modified to

$$\rho_{e,\text{mt}}(r) = \begin{cases} \rho_e(r) + \rho_e(2R_{\text{mt}} - r) + \rho_u & \text{for } r < R_{\text{mt}} , \\ 0 & \text{for } r > R_{\text{mt}} . \end{cases} \quad (2.76)$$

The constant ρ_u was introduced to ensure that the total charge is correct. The electrostatic potential is similarly modified to

$$V_{\text{st,mt}}(r) = V_{\text{st}}(r) + V_{\text{st}}(2R_{\text{mt}} - r), \quad (2.77)$$

and the total potential becomes

$$V_{\text{mt}}(r) = \begin{cases} V_{\text{st,mt}}(r) + V_{\text{ex}}(r) + V_{\text{cp}}(r) - iW_{\text{abs}}(r) & \text{for } r < R_{\text{mt}} , \\ V_{\text{st,mt}}(R_{\text{mt}}) + V_{\text{ex}}(R_{\text{mt}}) + V_{\text{cp}}(r) & \text{for } r > R_{\text{mt}} . \end{cases} \quad (2.78)$$

The muffin-tin potential remains rotationally symmetric. It is therefore unable to capture the full scope of solid-state effects, but it is a good first approximation.

The effect of these models on the differential scattering cross section can be seen in figure 2.9. This figure is for 100 eV electrons and silicon atoms. Note the logarithmic vertical axis: Mott scattering is predominantly forward-peaked.

The mean free paths and transport lengths can be found in figure 2.10. At ≥ 150 eV energies, the transport lengths of all models are very similar. This means that, averaged over a large number of scattering events, all models will produce similar behaviour. When the muffin-tin correction is enabled, the mean free path shifts up by a significant amount while the transport length is barely affected. This is because the Muffin-tin model's primary effect is to reduce the amount of forward scattering. At low $\lesssim 150$ eV energies, the models start to diverge as scattering becomes much more sensitive to the details of the interaction potential.

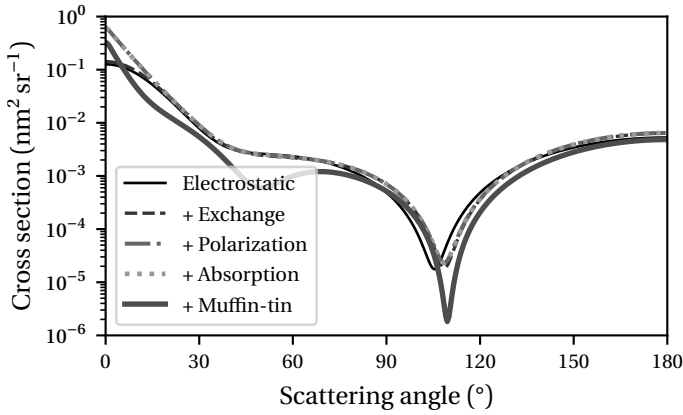


Figure 2.9: Differential Mott scattering cross sections at 100 eV for silicon, computed using ELSEPA. The models are applied in the same order as in the legend, meaning that the “polarization” curve includes both the exchange and correlation-polarization models.




Figure 2.10: The (a) mean free path and (b) transport length for silicon, calculated using ELSEPA. The various models discussed in this section are cumulatively enabled.

2.7 Phonon scattering

As we have seen, Mott scattering becomes sensitive to the details of the atomic potentials when the electron energy drops below $\lesssim 100$ eV. There is another issue to consider, however. The De Broglie wavelength of the electron is given by

$$\lambda = \frac{h}{p} \approx \frac{h}{\sqrt{2mE}}, \quad (2.79)$$

where h is Planck's constant ($2\pi\hbar$), m is the electron mass and E is its kinetic energy. Typical lattice constants in materials are on the order of 5 \AA . When E is on the order of 10 eV or less, the De Broglie wavelength is as large as interatomic distances and the electron cannot be considered to scatter on single atoms². If the material is crystalline, the energy eigenstates of electrons are Bloch waves. Bloch waves propagate freely through the crystal without scattering, instead only scattering when the crystal is imperfect. This happens when there are defects, but it can also happen due to lattice vibrations. The quanta of lattice vibrations are known as phonons.

When an atom is displaced by lattice vibrations, the potential changes locally. This gives rise to the electron–phonon interaction we will study in this section. There is an additional effect, called the Fröhlich interaction, which takes place in ionic lattices where displaced atoms have opposite charges. This gives rise to long-range electric fields. This second effect is only weak and will be ignored in this thesis. There are several good books covering the topic of electron–phonon scattering [73–75], only the main ideas will be presented here.

The theory uses the Born-Oppenheimer approximation. The atoms of the solid move very slowly compared to the electrons. The dynamics of the electrons and of the lattice can therefore be described separately, and the interaction is taken into account using perturbation theory. The scattering rate (per unit time) of electrons is given by Fermi's golden rule:

$$\Gamma = \frac{2\pi}{\hbar} |\mathcal{M}|^2 \delta(E_f - E_i), \quad (2.80)$$

where E_i and E_f are the energies of the initial and final states, respectively, and the matrix element

$$\mathcal{M} = \langle \psi_f | V | \psi_i \rangle. \quad (2.81)$$

²Diffraction is another elastic effect which cannot be considered to occur on single atoms. Diffraction also takes place at much higher energies if the electron's coherence length is greater than the interatomic distance. This section is not about diffraction.

V is the interaction potential. The final and initial wave functions are split up into an electronic and lattice part,

$$\psi_{\mathbf{k},n} = \phi_{\mathbf{k}} \times X_n, \quad (2.82)$$

where

$$\phi_{\mathbf{k}}(\mathbf{r}) = e^{i\mathbf{k}\cdot\mathbf{r}} u_{\mathbf{k}}(\mathbf{r}) \quad (2.83)$$

are the electron Bloch waves (u has the same periodicity as the crystal) and X_n represent the lattice. \mathbf{k} is the wave vector of the electron; n is the phonon occupation number.

2.7.1 Interaction potentials

There are several forms of the interaction potential. Two of them will now be discussed.

Deformation potential

The deformation potential interaction takes the following point of view. The change in potential energy is due to atoms being out of their equilibrium position. It is irrelevant what caused them to be out of equilibrium: a deformation of the entire material would cause the same effect as a displacement due to lattice vibrations. It is possible to measure this effect by stretching the material, which leads to a shift in the band edges.

The interaction potential is then given by

$$V = \sum_{i,j} \Xi_{i,j} u_{i,j}, \quad (2.84)$$

where $u_{i,j}$ are the components of the strain tensor and $\Xi_{i,j}$ are components of the deformation potential tensor. In crystals with spherical energy surfaces, the off-diagonal elements in the deformation potential tensor are zero and the diagonal elements are identical. In other crystals, this is still often a good approximation [76]. The interaction potential then reduces to

$$V = \Xi \nabla \cdot \mathbf{u}(\mathbf{r}). \quad (2.85)$$

Ξ is now a scalar, known as the deformation potential. $\mathbf{u}(\mathbf{r})$ is the displacement field, which describes by how much the atoms at position \mathbf{r} are displaced.

Rigid-ion approximation

In the rigid ion approximation, the full potential experienced by electrons is seen as the sum of individual atomic potentials. This potential stays attached to the ions as they are displaced.

The perturbation potential is equal to the potential difference between the situation where all atoms are in their equilibrium positions and them being in their out-of-equilibrium positions:

$$V = \mathcal{V}(x, \{\mathbf{X}\}) - \mathcal{V}(x, \{\mathbf{A}\}). \quad (2.86)$$

Here, x represents the electron positions, $\{\mathbf{X}\}$ the set of all true atom positions and $\{\mathbf{A}\}$ their equilibrium positions. In the rigid-ion approximation, \mathcal{V} is given by

$$\mathcal{V}(r, \{\mathbf{R}\}) = \sum_{\mathbf{n}, \alpha} v^\alpha(\mathbf{r} - \mathbf{R}_\mathbf{n}^\alpha), \quad (2.87)$$

where the sum over \mathbf{n} is over the unit cells in the lattice and α labels the atom within the cell. v is the potential that each atom carries rigidly with it.

Performing a Taylor expansion for small displacements $\mathbf{U}_\mathbf{n}^\alpha$, it can be shown that the interaction potential reduces to

$$V = - \sum_{\mathbf{n}, \alpha} \mathbf{U}_\mathbf{n}^\alpha \cdot \nabla v^\alpha(\mathbf{r} - \mathbf{A}_\mathbf{n}^\alpha), \quad (2.88)$$

where $\mathbf{A}_\mathbf{n}^\alpha$ is the equilibrium position of the atom.

2.7.2 Matrix element

With appropriate expressions for the lattice wave functions and the atomic displacements, the matrix element can be evaluated. It turns out that the only nonzero matrix elements satisfy

$$\mathcal{M}_{\mathbf{k}'n', \mathbf{k}n}^\pm = \sqrt{\frac{\hbar}{2M\omega_{\mathbf{q}j}}} \sqrt{n_{\mathbf{q}j} + \frac{1}{2} \mp \frac{1}{2}} I^\pm, \quad (2.89)$$

$$I^\pm = \int_{\Omega_0} d^3\mathbf{r} u_{\mathbf{k}\pm\mathbf{q}}^* (\hat{\mathbf{e}}_{\mathbf{q}j} \cdot \mathbf{f}^\pm) u_{\mathbf{k}}. \quad (2.90)$$

M is the total mass of the solid; $\omega_{\mathbf{q}j}$ is the frequency of a phonon with wave vector \mathbf{q} and polarization j (there are two transverse and one longitudinal polarizations); $\hat{\mathbf{e}}_{\mathbf{q}j}$ is that phonon's polarization vector; $u_{\mathbf{k}}$ is the periodic part of the electron's Bloch wave; and the integration domain Ω_0 is the elementary unit

cell. The vector \mathbf{f}^\pm depends on the choice of interaction potential and will be discussed later. $n_{\mathbf{q}j}$ is the phonon occupation number, which is given by the Bose-Einstein distribution,

$$n_{\mathbf{q}j} = \frac{1}{\exp(\hbar\omega_{\mathbf{q}j}/k_B T) - 1}. \quad (2.91)$$

The upper and lower signs in equation (2.89) refer to phonon emission and absorption, respectively. Matrix elements involving more than one phonon simultaneously are zero. The selection rules are $\mathbf{k}' = \mathbf{k} - \mathbf{q} + \mathbf{g}$, $E' = E - \hbar\omega_{\mathbf{q}j}$ for phonon emission, and $\mathbf{k}' = \mathbf{k} + \mathbf{q} + \mathbf{g}$, $E' = E + \hbar\omega_{\mathbf{q}j}$ for phonon absorption, respectively. \mathbf{g} is a reciprocal lattice vector.

The reciprocal lattice vector \mathbf{g} must be chosen such that $\mathbf{k}' - \mathbf{k}$ falls inside the reduced Brillouin zone. Processes where $\mathbf{g} = 0$, which is always the case when the electron has very low energy, are called normal processes or N-processes. When $\mathbf{g} \neq 0$, they are called umklapp processes or U-processes.

Deformation potential

For the deformation potential interaction, \mathbf{f}^\pm is

$$\mathbf{f}^\pm = \pm i\Xi\mathbf{q}. \quad (2.92)$$

Assuming an isotropic medium, the coupling factor (2.90) reduces to

$$I^\pm = \pm i\Xi(\hat{\mathbf{e}}_{\mathbf{q}j} \cdot \mathbf{q}). \quad (2.93)$$

The factor $\hat{\mathbf{e}} \cdot \mathbf{q}$ implies that only longitudinal modes contribute, since the dot product is zero for transverse modes. This statement is not generally true, though. It depends on the assumption that the medium is perfectly isotropic, which is not true for a real crystal. In general, both longitudinal and transverse modes should be taken into account. This can be done by replacing $\hat{\mathbf{e}} \cdot \mathbf{q} \rightarrow q$ in equation (2.93), and introducing separate longitudinal and transverse deformation potentials. There are then six electron-phonon scattering rates: there is phonon emission and absorption for each of the three phonon branches (one longitudinal and two transverse).

Rigid-ion approximation

For the rigid-ion approximation, \mathbf{f}^\pm is given by

$$\mathbf{f}^\pm = - \sum_{\mathbf{n}} e^{\pm i\mathbf{q} \cdot (\mathbf{n} - \mathbf{r})} \nabla v(\mathbf{r} - \mathbf{n}). \quad (2.94)$$

Taking plane waves for the Bloch functions $u_{\mathbf{k}}$, the coupling factor (2.90) can be reduced to

$$I^{\pm} = \mp \frac{i}{\Omega_0} (\hat{\mathbf{e}}_{\mathbf{q}j} \cdot \mathbf{q}) \int_{\Omega} d^3\mathbf{r} e^{\mp i(\mathbf{q}\cdot\mathbf{r})} v(\mathbf{r}). \quad (2.95)$$

Here, the integral over Ω is over all space, not the unit cell. The coupling factor is therefore proportional to the Fourier transform of the atomic potential $v(\mathbf{r})$.

We are now ready to plug in a shape for the atomic potential $v(\mathbf{r})$. An obvious choice is the screened Coulomb potential already mentioned in equation (2.74). Use of this potential for phonon scattering was suggested by Bradford and Woolf [77], and it was also used by Schreiber and Fitting [78–80]. The coupling factor then reads

$$I^{\pm} = \pm i \Xi (\hat{\mathbf{e}}_{\mathbf{q}j} \cdot \mathbf{q}) \frac{a^2}{q^2 + a^2}. \quad (2.96)$$

An overall multiplicative constant, $Z e^2 / (\epsilon_0 \Omega_0 a^2)$, was replaced by the deformation potential Ξ . This follows the paper of Bradford and Woolf [77], who recognised that the deformation potential interaction of equation (2.93) is recovered in the limit $q \rightarrow 0$. a is the screening parameter of equation (2.74). Bradford and Woolf estimate its value as $a^2 = Z e^2 / (\epsilon_0 \Omega_0 \Xi)$. Schreiber and Fitting [78–80] leave it as a fitting parameter. Compared to the plain deformation potential interaction, the screened Coulomb potential suppresses phonon scattering at high q .

Another popular choice for the atomic potential is to take the ion as a sphere with radius r_s in which its charge is homogeneously distributed. This approach is taken, amongst others, by Pop *et al.* [81] and Akkerman and Murat [82]. In that case,

$$I^{\pm} = \pm i \Xi (\hat{\mathbf{e}}_{\mathbf{q}j} \cdot \mathbf{q}) 3 \frac{\sin(qr_s) - qr_s \cos(qr_s)}{(qr_s)^3}. \quad (2.97)$$

An overall constant was again recognised as the deformation potential Ξ . As was the case for the screened Coulomb potential, the factor at the end is unity when $q = 0$ and gradually decreases as q increases.

Figure 2.11 illustrates these model potentials for silicon. This figure shows $\tilde{V}(q) = \frac{1}{\Omega_0} \int d^3r e^{i\mathbf{q}\cdot\mathbf{r}} v(r)$, where Ω_0 is the unit cell volume and $v(r)$ is the atomic potential. Two potentials mentioned above are shown: the screened Coulomb potential used by Schreiber and Fitting and the “uniformly charged ion” potential used by Pop *et al.*. We show another potential based on numerically calculated charge densities of the electron cloud, for an isolated cold neutral atom. These are the same Dirac-Fock densities used in the ELSEPA package [72].




Figure 2.11: Three different atomic potentials for silicon. Parameters used are $\Xi = 6.4 \text{ eV}$ [81], $a = 52 \text{ nm}^{-1}$ (consistent with [10]), $r_s = 2.1 \text{ \AA}$ [81].

The screened Coulomb and uniform ion potentials are equal to the deformation potential at $q = 0$. The screened Coulomb potential decays much more slowly with increasing q . The result is that large-angle scattering (which involves large momentum transfers q) is much more likely. The mean free path will therefore be shorter and the angular distribution will be more uniform, leading to a much smaller transport length. The Dirac-Fock potential is similar in shape to the screened Coulomb potential, but it is much lower. The angular distribution will therefore be similarly uniform, but the mean free path will be significantly larger because the interaction is less strong.

2.7.3 Scattering rate

We have already mentioned Fermi's golden rule, equation (2.80). We are interested in the transition rate from a known initial state with wave vector \mathbf{k} , into any final state \mathbf{k}' . Hence we must sum (or integrate) over all final states:

$$\Gamma^\pm = \frac{\Omega}{(2\pi)^3} \int d^3\mathbf{k}' \frac{2\pi}{\hbar} |\mathcal{M}^\pm|^2 \delta(E_{\mathbf{k}'} - E_{\mathbf{k}} \mp \omega_{\mathbf{q}}). \quad (2.98)$$

The integral over the magnitude of \mathbf{k}' can be solved using the delta function. Inserting equation (2.89) for the matrix element, the scattering rate becomes

$$\Gamma^\pm = \frac{m^*}{8\pi^2 \hbar^2 \rho} \int d\cos\theta \int d\phi \frac{k'}{\omega_{\mathbf{q}}} |I^\pm|^2 \left(n_{\mathbf{q}j} + \frac{1}{2} \mp \frac{1}{2} \right). \quad (2.99)$$

ρ is the density of the material, m^* is the electron effective mass. θ and ϕ are, respectively, the polar and azimuthal angles between \mathbf{k} and \mathbf{k}' (that is, the scattering angle of the electron). Note that there are two relations between \mathbf{k}' , \mathbf{k} , and \mathbf{q} : $\mathbf{k}' = \mathbf{k} \pm \mathbf{q} + \mathbf{g}$ and $E' = E \pm \omega_{\mathbf{q}j}$. Therefore, k' , $\omega_{\mathbf{q}j}$, I^\pm and $n_{\mathbf{q}j}$ are implicitly a function of θ . The angular distribution of the electron in a scattering event is given by the integrand of equation (2.99). The mean free path can be obtained from $\lambda = \sqrt{2E/m^*}/\Gamma$.

The equation quoted by Pop *et al.* [81] and Akkerman and Murat [82], which is only used for N-processes, can be obtained by inserting equation (2.97) into (2.99).

Schreiber and Fitting [78–80] are interested in a large range of electron energies, so U-processes are possible as well as N-processes. Remember that $\mathbf{k}' = \mathbf{k} \mp \mathbf{q} + \mathbf{g}$, where \mathbf{g} is a reciprocal lattice vector. One speaks of an N-process when $\mathbf{g} = 0$ and of an U-process otherwise. When $k < k_{BZ}/2$, where k_{BZ} is the size of the Brillouin zone, only N-processes are possible. U-processes dominate when $k > k_{BZ}$. In that limit, Sparks *et al.* [83] propose taking “suitable weighted averages of zone-boundary phonon frequencies, and an average over the directions of the phonon eigenvector”. Schreiber and Fitting interpret this as replacing $\omega_{\mathbf{q}j} \rightarrow \omega_{BZ}$ in equation (2.99). Schreiber and Fitting make a few more approximations:

- Assuming the phonon energy is much less than the electron energy, $k' \approx k$ and $q^2 \approx 2k^2(1 - \cos\theta)$
- For N-processes, there is a linear dispersion between the phonon wave vector and its energy: $\omega_{\mathbf{q}j} = c_s q$. c_s is the speed of sound.
- For N-processes, the Bose-Einstein distribution is approximated as

$$2n + 1 \approx \frac{2k_B T}{\hbar c_s q}. \quad (2.100)$$

- They define a quantity

$$A = \frac{\hbar^2 a^2}{8m^*}, \quad (2.101)$$

where a is the screening parameter from equation (2.96).

With the coupling factor from equation (2.96), the total (for emission + ab-

sorption) differential scattering rate finally becomes³

$$\frac{d\Gamma}{d\Omega} = \begin{cases} \frac{k_B T \Xi^2}{4\hbar c_s^2 \rho} D(E) \frac{1}{\left[1 + \frac{E(1-\cos\theta)}{2A}\right]^2} & \text{for } k < k_{BZ}/2, \\ \frac{m^* \Xi^2 E}{2\rho \hbar^2} D(E) \frac{2n_{BZ} + 1}{\omega_{BZ}} \frac{1 - \cos\theta}{\left[1 + \frac{E(1-\cos\theta)}{2A}\right]^2} & \text{for } k > k_{BZ}. \end{cases} \quad (2.102)$$

In between these limiting cases, Schreiber and Fitting employ a linear interpolation between the two mechanisms.

For $k < k_{BZ}/2$, the angular distribution follows a trend familiar from other kinds of scattering events: scattering mostly takes place in the forward direction at high electron energies and it becomes more isotropic towards lower energies. For $k > k_{BZ}$, however, forward scattering is suppressed. The differential cross section has a maximum at nonzero scattering angles. In addition, Schreiber and Fitting had trouble matching simulation to experiment. They therefore chose to use the angular distribution for $k < k_{BZ}/2$ at all electron energies [78, page 72]. This is not clear from electronically available publications [79, 80]. As a result, several others [10, 12, 13] used the as-derived differential scattering rates instead.

All angular distributions are illustrated in figure 2.12. The material is silicon. This figure shows the angular distribution of Akkerman and Murat [82] (using the “uniform ion” potential from figure 2.11), that of Schreiber and Fitting [78–80], and one derived using the numerical Dirac-Fock potential. Each of these angular distributions were evaluated at an electron energy of 3 eV, corresponding to N-processes. In addition, we show the angular distribution of Schreiber and Fitting evaluated at 50 eV energy, corresponding to U-processes. All parameters were taken from reference [82]. The screening parameter for the Schreiber-Fitting model is $A = 5E_{BZ}$.

The corresponding mean free paths and transport lengths are shown in figure 2.13. As only the Schreiber-Fitting model describes U-processes, only that model is used at high electron energies. At low energies, the respective mean free paths and transport lengths are very similar. This is because of the relatively uniform angular distribution seen in figure 2.12. At high energies, the transport length of Schreiber and Fitting is slightly shorter than the mean free path. That is because the model favours backscattering. Most strikingly, the Schreiber-Fitting mean free path is significantly shorter than the other two.

³Ref. [80] misses a factor \hbar in the denominator of the expression for $k > k_{BZ}$.

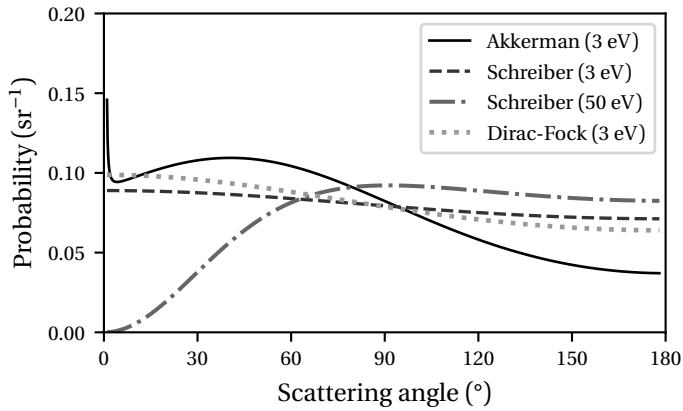


Figure 2.12: Angular distributions for several electron–acoustic phonon scattering models discussed in the main text. The material is silicon.




Figure 2.13: The (a) mean free paths and (b) transport lengths for several electron–acoustic phonon scattering models discussed in the main text. The material is silicon, energies are referenced with respect to the bottom of the conduction band.

Another thing to consider is the value of the deformation potential. Fischetti and Laux [84] give a good summary of reported values. They report values for the shear (Ξ_u) and dilation (Ξ_d) deformation potentials, which, after averaging over all directions, are related to longitudinal and transverse deformation potentials by [81]

$$\Xi_{LA} = \sqrt{\frac{\pi}{2}} \sqrt{\Xi_d^2 + \Xi_d \Xi_u + \frac{3}{8} \Xi_u^2}, \quad (2.103)$$

$$\Xi_{TA} = \frac{\sqrt{\pi}}{4} \Xi_u. \quad (2.104)$$

Fischetti and Laux give values for Ξ_u ranging from 7.5–10.5 eV, and Ξ_d in the range of -11.7 – 1.1 eV, all for silicon. Ξ_{LA} is then between 3.3 eV (when $\Xi_d = -3.75$ eV, $\Xi_u = 7.5$ eV) and 9.2 eV (when $\Xi_d = 1.1$ eV, $\Xi_u = 10.5$ eV). Ξ_{TA} is in the 3.3–4.7 eV range. Because the electron–phonon scattering rate depends quadratically on the deformation potential, this range of deformation potentials translates to a substantial uncertainty in the scattering rate.

The mean free path near the Fermi level can also be estimated from the electron mobility, or, in case of metals, conductivity. The scattering rate can be obtained using

$$\Gamma = \frac{e}{m^* \mu_e} = \frac{e^2 n}{m^* \sigma}. \quad (2.105)$$

μ_e is the electron mobility, σ is conductivity, e is the elementary charge and n is the number density of conductance electrons. This model assumes that the electron scatters isotropically in every interaction, so the corresponding path length $\lambda = \sqrt{2E/m^*}/\Gamma$ should be interpreted as the transport length. Due to the isotropic nature of electron–acoustic phonon scattering at low energy, this is close to the mean free path. When this model is used for metals, E is the Fermi energy with respect to the bottom of the band. When used for semiconductors or insulators, E is the energy of the conduction electrons above the bottom of the conduction band, which can be estimated as $\frac{3}{2} k_B T \approx 39$ meV at room temperature. For Si, $\mu_e = 1400 \text{ cm}^2 \text{ V}^{-1} \text{ s}^{-1}$ [85] and $m^* = 0.26 m_e$ [86], which leads to $\lambda = 47$ nm. As the resulting scattering rate includes scattering due to defects and optical phonons as well as acoustic phonons, this value represents a lower bound.

2.7.4 Conclusion

With all the physics laid out in the previous sections, it is now time to choose a model. The calculation of Akkerman and Murat [82] was specifically tailored to silicon, and is very difficult to repeat for other materials. We have shown a similar rigid-ion approximation calculation using numerical Dirac-Fock potentials. This approach ignores the deformation of the atomic potentials when there are lattice vibrations, and cannot be calibrated to include an experimental deformation potential. This approach is very rarely used in literature. The method of Schreiber and Fitting [78–80] is the most general one described here, and also works at energies large enough that U-processes dominate. That makes it the most convenient for implementation in our simulator.

One problem that exists with this model is that the deformation potential is often not readily available, or there is a very large variation in literature values. The electron mobility (for semiconductors) or resistivity (for metals) is much more accurately known for many metals. We will therefore calibrate the deformation potential by the electron mobility or resistivity. We require that the electron–acoustic phonon scattering rate of equation (2.102) at the bottom of the conduction band (for semiconductors) or at the Fermi energy (for metals) equals that of equation (2.105).

2.8 Summary

We will now summarize our choice of models in this section.

We describe inelastic scattering, in which the electron loses energy, by the dielectric function model of section 2.4. The mean free path is given by equation (2.34) and the only material parameter is the dielectric function, $\text{Im}[-1/\epsilon(q, \omega)]$. We find this by combining measured optical data $\text{Im}[-1/\epsilon(0, \omega)]$ with the full Penn algorithm [42].

The dielectric function model describes how an electron loses energy and momentum, but not where that energy and momentum goes. The models used for this are described in section 2.5. A secondary electron is taken from the valence band in most cases, either through direct excitation or indirectly through the creation of a plasmon. Ionization of target atoms' inner shells is also possible.

We choose a hybrid model for elastic scattering. At high energies of several hundred eV and more, it is known empirically that Mott scattering (section 2.6) dominates. We use Mott scattering there exclusively. Meanwhile, it is known

that electron–phonon scattering (section 2.7) describes conduction phenomena very well. Therefore, at energies around the Fermi level, this type of scattering should dominate. As Mott scattering does not naturally “turn off” at low energies, we must turn it off artificially. Similarly, the electron–phonon scattering mechanism from Schreiber and Fitting [78–80] does not naturally turn off at higher energies, so we must also do that artificially. We use a mechanism originally proposed by Kieft and Bosch [10] and subsequently copied by Verduin [12] and Theulings [13]. We use phonon scattering at < 100 eV and Mott scattering at > 200 eV. In between these boundaries, we interpolate between the two. These limits were chosen arbitrarily. Other options will be explored in section 6.3.

The final type of scattering is material boundary crossing. An electron must cross a potential barrier when escaping a material. The electron may be reflected on this barrier, or it may change direction if it crosses the barrier. This was described in section 2.3.

Chapter 3

Monte Carlo simulator

Having discussed the physics of electron-matter interaction in chapter 2, we now turn to the problem of implementing these models in a computer simulation. We will use the Monte Carlo method. Section 3.1 introduces the general principle of the simulation. The simulator can run both on graphics processing units (GPUs) and on central processing units (CPUs). The technicalities of how this is achieved are discussed in section 3.2. We will then gain more insight in the electron scattering process in sections 3.3 and 3.4. Simulation results will then be compared to experimental electron yields in section 3.5. Finally, we will perform a sensitivity analysis for scattering processes in section 3.6.

*The simulator described in this chapter has been published as an open-source software package (<https://nebula-simulator.github.io>). It was also published as L. van Kessel, C.W. Hagen, “Nebula: Monte Carlo simulator of electron-matter interaction”, *SoftwareX* **12**, 100605 (2020) [87].*

3.1 Introduction

From a physical perspective, we make the following main assumptions. An electron is treated as a classical point particle, scattering through the volume of the material in discrete events. The electron is treated as if it is in free flight between such events. We also assume that all electrons can be treated independently.

As described in section 2.1, the distance between two scattering events is a random variable. What happens in a scattering event—for example by what amount an electron is deflected—is also random. The general idea of our Monte Carlo simulator is to “act out” one specific realisation of what might happen to an electron once it reaches the sample. The simulator literally traces the path

that the electron takes, as well as that of all the secondary electrons that may be produced along the way. The output of the simulator is, therefore, also random, and starting it again will yield a different result. We need to provide the simulator with appropriate probability distributions for electron scattering. The physical origin of these has already been described in chapter 2. These distributions are calculated before the simulation starts. The simulator is equipped with a random number generator which allows us to sample random situations according to the proper probability distributions.

The simulator will distinguish three basic types of scattering. The first is inelastic scattering, where the electron loses energy and may create a secondary electron. We use the Penn algorithm (section 2.4), to determine the differential scattering cross sections and the method described in section 2.5 for secondary electron generation. The second kind of scattering is elastic scattering, for which we use Mott scattering (section 2.6) at > 200 eV energies, phonon scattering at < 100 eV, and interpolation in between. Finally, on the boundary between two materials, the electron may reflect or refract as described in section 2.3.

We will now walk through one iteration of the simulator. At the beginning of an iteration, an electron has a known position, direction and energy. If the electron is in vacuum, it cannot scatter elastically or inelastically, and it is moved to the next intersection with a material boundary in a straight line. If the electron is in a material, the elastic and inelastic scattering processes are probed for random free path lengths. The exact distance of travel to the next event is a random number. Its probability distribution is typically given by the exponential Lambert-Beer law, with the mean free path depending on the electron's energy. The nearest of the three events (elastic, inelastic or boundary crossing) is chosen. The electron is moved in a straight line to the next event, and the event is performed. The outcome of the event, which may *e.g.* involve a deflection or energy loss, is generally random. A material boundary crossing may involve reflection, refraction and/or an energy gain due to the difference in inner potential between the materials. A scattering event in the bulk of the material may involve a deflection of the electron, energy loss, and the creation of a secondary electron. If a secondary electron is created, it is added to the list of active electrons and processed in the same way as the primary electron. The procedure is repeated until an electron reaches a detector, or its energy drops below a certain level. This cutoff energy level is usually set such that we stop tracking an electron when its energy is insufficient to leave the material.

At this point, we should highlight the “hybrid” nature of these simulations.

From the simulator's point of view, the electron is a classical (non-quantum) point particle that scatters in discrete events. The physical description of the scattering mechanisms themselves, though, can be based in quantum mechanics. As an example, consider Mott scattering described in section 2.6. This is a fully quantum mechanical description of an electron scattering on a single isolated atom. The result of the Mott scattering calculation is a differential scattering cross section $d\sigma / d\Omega$, which is converted to a differential inverse mean free path $d\lambda^{-1} / d\Omega = n d\sigma / d\Omega$, where n is the number of atoms per unit volume. In this step, we have implicitly switched to a description in which the electron is a classical point particle, not a wave packet, and in which the material looks like a homogeneous "soup" of atoms. Situations in which both the wave properties of electrons and the relative positioning of the atoms are important, such as diffraction experiments, cannot be described because of this step. In some situations, it is possible to derive a single scattering cross section that takes the collective positioning of atoms into account. This is what we did for electron-acoustic phonon scattering in section 2.7. Ideally, though, one would perform a calculation in which the electron is treated as a wave entirely.

Charging effects cannot be described due to the assumption that all electrons are independent, We also ignore electric and magnetic fields in this simulation, so the electrons always move in straight lines between scattering events.

3.2 Implementation

The main simulator, which we called Nebula, takes three types of input:

- *The geometry*, which is provided as a list of triangles representing the interfaces between materials. The triangles know which materials are on either side. The shape of the detector is also part of the geometry, and is implemented as a special detector material.
- *The primary electrons*, which are provided as a list of starting positions, directions, energies and tags. The tags, which are two integer values specified by the user, have no physical meaning in the simulator. Their values do not change throughout the electron's lifetime, and they are passed to all secondary electrons in the cascade. Their intended use is to represent the pixel in the SEM image that the electron belongs to.
- *Material file(s)*, one for each material in the sample. These store detailed scattering cross sections, ready for use in the simulation. For example, in




Figure 3.1: Illustration of the data flow in the software. Nebula is the main simulator, cstool compiles scattering cross sections for a material.

the case of Mott scattering, there are two tables: the mean free path as function of energy; and for each energy, the probability density function¹ of the angular scattering distribution.

The output is a list of electrons that hit a detector material, in the same format as the primary electrons. For each detected electron, then, we store its position, direction, energy and tags at the moment of the detection event.

The material files are compiled with a separate tool, called cstool (short for “cross section tool”). It takes as input a human-friendly text file with material parameters, such as the elemental composition and dielectric function. The parameters used in this thesis are given in appendix A. It runs the ELSEPA package [72] to compute Mott scattering cross sections, extracts necessary data from the LLNL Evaluated Electron Data Library [59], and contains its own implementation of the Penn algorithm [42]. The cross section tool is intended to be run only once for each material of interest.

The entire flowchart is depicted in figure 3.1.

¹Technically, we store the inverse of the cumulative distribution function because this makes sampling random numbers from this distribution at run-time very fast.

3.2.1 GPU simulation

Simulations of this type are notoriously slow. It often takes hours to simulate a full scanning electron microscope image. In an attempt to maintain a reasonable speed, some authors have resorted to semi-empirical models [10, 88, 89].

Verduin *et al.* [90] demonstrated Monte Carlo simulations running on graphics processing units (GPUs). Specifically, they used Nvidia's CUDA architecture. A regular central processing unit (CPU) in a computer has only a few compute cores, typically between 2 for entry-level consumer hardware and 24 for high-end units in compute clusters. By contrast, a good GPU has thousands of cores. The catch is that, while the cores in a CPU can be programmed as if they are fully independent, in a GPU they cannot. On a GPU, many cores need to execute the same program simultaneously. In addition, they operate in groups known as *warps* which execute instructions in lockstep². If instruction divergence occurs within a warp, for example, due to an if-then block or loops of unequal length, some threads in the warp are put on hold instead of advancing to the next instruction. This should be avoided because it represents a loss of performance. Apart from warp divergence, there are other factors to consider, including memory latency, memory bandwidth and register pressure. Today's processors (both CPUs and GPUs) are often so fast that they can execute many instructions per byte of data transferred from memory. Consequently, many programs are limited by memory bandwidth. GPUs have an advantage over CPUs here because their memory throughput is much higher.

The main difficulty in running an electron scattering simulator on the GPU lies in the fact that electrons may generate secondary electrons. It is not possible to dynamically allocate new memory to store these secondary electrons while the GPU program is running. Our simulation will therefore have a fixed maximum number of electron slots. We call this maximum number the electron capacity of the simulator. Some of these slots must be kept "free" so they can hold secondary electrons generated by inelastically scattering electrons. In addition, we must avoid situations in which two threads attempt to create a new secondary electron in the same memory location simultaneously. The solution of Verduin *et al.* [90] is the following. They assume that only one secondary electron may be generated in each inelastic event³. An iteration of the simulator is split into three steps. In the first step, which we will call the "initialization" step,

²This is not always truly the case, depending on the hardware. The programmer cannot rely on threads operating in lockstep.

³The method easily generalises to higher numbers if there is an upper bound that is a priori known.

each electron decides what its next event type is: inelastic scattering, elastic scattering or crossing a material boundary. Then, the electrons are sorted by type: those that will scatter inelastically go first, then those that will scatter elastically or cross a material boundary, and the empty slots go at the end. In the third and final step, the event for each respective electron is performed. If an electron with index i scatters inelastically, it will attempt to create its secondary electron in slot $N - i - 1$, where N is the electron capacity and the -1 is because our indexing starts from 0. This means that every inelastically scattering electron can find a unique slot in which to place a secondary. If that slot is already occupied, the event is not performed; the electron is put on hold until the next iteration.

The sorting step also solves much of the warp divergence problem. The code for scattering inelastically is very different from scattering elastically. Because electrons are now ordered by event type, most warps only have to perform one type of scattering event.

Geometry intersections

As mentioned before, the geometry is provided as a list of triangles specifying interfaces between volumes. This is quite conventional in many ray-tracing applications, though alternatives have also been used for electron microscopy purposes [10, 15, 91]. We use the well-known Möller-Trumbore algorithm for ray-triangle intersections [92] to determine whether an electron (travelling in a straight line) intersects a triangle.

Iterating through a list with potentially millions of triangles is an expensive operation. Fortunately, there are methods to sort the triangles in such a way that not all triangles need to be tested. Many such data structures have been suggested, such as uniform grids, octrees, and bounding volume hierarchy trees [93]. The basic idea behind all of these is to subdivide all space into cells. It is fairly easy to determine which cell an electron is in. We only need to perform ray-triangle intersection tests against the triangles occupying the same cell as the electron. If there are no intersections, the electron moves on to the next cell and the process is repeated. The procedure is stopped when a triangle intersection is found, or when it is determined that a regular scattering event in the bulk of the material comes first. The differences between the various types of data structures lie in the way the cells are chosen.

We use the same octree data structure as Verduin [12]. An octree divides space into eight equally-sized cuboid octants. Each of these may be recursively split into more octants, forming a tree-like structure. The triangles are assigned

to the “leaf” cells. If there are more than 16 triangles in a cell, the cell is subdivided into octants and the triangles are distributed among the new leaf cells. This keeps the number of triangles per cell low, and concentrates more cells in areas with dense geometries. If an electron is far away from dense geometry, it is usually in a large cell and is likely to undergo an elastic or inelastic scattering event before it has to check the triangles in the next cell.

A triangle may span more than one leaf cell, in which case it becomes a member of all cells it overlaps with. This means that a specific triangle is tested more than once if the electron moves between two of those cells. There are alternatives to the octree which avoid this, such as the bounding volume hierarchy tree, but they have irregularly shaped cells making it more expensive to find the cell the electron is in. This is ultimately a tradeoff, and it is currently unclear which option is best for the purposes of our electron microscopy simulator.

Running method

In a typical use case, the GPU has an insufficient amount of memory to hold all primary electrons and their subsequent cascades of secondary electrons simultaneously. Therefore, new primary electrons must be sent to the GPU in batches.

The size of a typical cascade as function of time is shown in figure 3.2. In the first few iterations, the primary electron has a lot of energy, so it can create many secondary electrons. Some of these may also have enough energy to create a third generation. The number of electrons per primary therefore rapidly increases. Soon, though, many electrons do not create secondary electrons. Instead, they are instead stopped by inelastic events. The number of active electrons reaches a maximum and then there is a long tail in which their number slowly decreases.

The long tail makes it impractical to wait for the previous batch of primary electrons to finish simulating before sending the next batch to the GPU. That would mean the GPU is under-utilized for most of the time. It is more natural to add the next batch of electrons when the previous batch has just reached its maximum size and starts to decay. Before the simulation starts, we do a “prescan” which measures the curve of figure 3.2. This gives us the iteration number in which the cascade size is maximal. The prescan also enables us to choose the batch size, that is, the number of primary electrons to be added in each batch. The expected number of active electrons during the simulation can be estimated ahead of time by repeatedly superimposing the prescan curve of figure 3.2 multiplied by the batch size. The result will tend to a steady state, which




Figure 3.2: Size of the electron cascade, measured as number of running electrons per primary electron, as function of time. This was measured for an infinitely large flat silicon sample, for different energies of the primary electron beam.

we aim to be 90% of the total electron capacity.

We have discussed how new input is provided to the simulation. We still need to mention how the output mechanism works. When an electron reaches a detector, we want to store its information. Instead of copying it to a separate buffer immediately, we keep it in the simulation memory. An electron slot can either be active (scattering inelastically, elastically or crossing a boundary); empty, or detected. Detected electrons do not participate in the simulation, just like empty electron slots, but they cannot be filled by secondary electrons either. When a new batch of primary electrons is added to the simulation, we copy the entire simulation memory over to the CPU and detected electron slots are marked as empty.

We should mention one final detail. The original implementation of Verduin [12] interrupts the simulation for a long time to add new electrons and clear the detected ones. When a set of iterations finishes, it copies the entire simulation memory from the GPU to the CPU. The CPU then outputs the detected electrons, marks their slots as cleared, and adds the new primary electrons. Then the entire simulation memory is copied back to the GPU. The problem with this is that copies between the CPU and GPU are slow, and the GPU has to wait unnecessarily while the CPU writes its output to disk. We now use a more efficient scheme, which uses CUDA's ability to perform CPU-GPU copies asynchronously while the simulation is running. While the GPU is working, we

already copy the next batch of primary electrons to a buffer on the GPU. When the time comes to add this batch to the simulation, we do two things on the GPU. First, we copy the entire simulation memory to a second buffer on the GPU. This second buffer will provide the detected electrons to the CPU later. Second, we perform a sorting step and run a simple GPU kernel to add the new electrons to free electron slots while setting detected electron slots to empty. Then the simulation continues. While the simulation runs again, we copy the output buffer to the CPU, which writes the output to disk. The CPU sends new primary electrons to the input buffer on the GPU. This scheme only interrupts the simulation to copy data within GPU memory itself. Such copies are much faster than copies to the CPU. It saves about 20% in the overall simulation time.

3.2.2 CPU simulation

The GPU is a great tool for fast simulations. However, it is not very flexible, and it is difficult to efficiently extract information that the simulator was not designed to produce. For example, we might want to produce a three-dimensional map of energy deposited by the electrons. This could be interesting for electron-beam lithography applications, where a common model is that the solubility of resist is closely related to the deposited energy. Alternatively, we might be interested in a map of positive and negative charges. While it is certainly possible to investigate both these situations on the GPU, it requires a lot of additional bookkeeping. It is easier to run such simulations on the CPU, where it is much easier to output arbitrary pieces of data.

Fortunately, this is possible thanks to the syntactic similarity between CUDA code and ordinary C++. This means that much of the code, including the physical models and the octree implementation, can be shared and compiled for either CPU or GPU as needed. There are only two areas where special care is needed.

The first is electron management. As mentioned, the GPU simulator has a fixed electron capacity, which necessitates special memory management and a sorting step. Additionally, an elaborate setup is needed to add new electrons to the simulation and get the detected electrons out. None of this is necessary on the CPU, where new and secondary electrons can be added to a resizable stack as desired, and there is no need to process thousands of electrons simultaneously. All this management code is therefore moved into “drivers”. These are responsible for accepting new electrons to be simulated, adding secondary electrons to the simulation, and extracting detected electrons. They also query




Figure 3.3: Illustration of the CPU-GPU split. Drivers are responsible for managing the electrons in the simulation. Physics code is shared between CPU and GPU versions. A `gpu` template parameter is passed to the data storage classes, which copy data to the GPU if necessary.

the physics models and geometry for free path lengths, move electrons to their scattering positions, and call the physics models to perform events. The GPU driver is also responsible for the sorting step.

The second is data storage for the physical models. While the physics code itself is device-independent, it is often necessary to store some data, such as a table of mean free paths as function of energy. In case of a GPU simulation, this data must be copied to the GPU beforehand. The solution is to have special data storage classes, which carry a `gpu` template parameter. If this template parameter is true, these classes copy their data to the GPU. The CPU and GPU drivers, which own the physics classes, pass the correct template parameter to the physics classes. The physics classes don't need this parameter directly, but they pass it on to the data storage classes. The organisation is illustrated in figure 3.3.

The CPU version comes with additional advantages, such as the ability to add an arbitrary number of scattering mechanisms. The GPU version currently only supports two, though this restriction can be relaxed quite easily. The CPU version makes it easier to add *e.g.* a trapping mechanism in addition to inelastic and elastic scattering. It is also possible to create an arbitrary number of secondary electrons per scattering event in the CPU version. These properties are nice to have for experimentation with physical models.

The CPU code can be easily parallelised across multiple cores. Every core takes a single primary electron, and maintains a stack of secondaries created in the cascade. It simulates the cascade until it dies out, then requests a new primary electron from the pool of primaries. It is also possible for the GPU simulator to run CPU simulations in parallel to the GPU simulation on unused cores,




Figure 3.4: (a) throughput of a simulation on the CPU and GPU, as function of primary electron energy. (b) speedup of a GPU simulation compared to the CPU. Our CPU simulations were done on two Intel Xeon E5-2683 v4 CPUs, the GPU simulations were performed on a single Nvidia Geforce GTX1080.

but this has not been implemented yet. At the moment, there are only separate CPU-only and GPU-only simulators.

3.2.3 Speed comparison

It is now interesting to compare the speed of the GPU simulation to the CPU. The machine we have at our disposal has two CPUs, each an Intel Xeon E5-2683 v4 with 16 cores. With hyperthreading enabled, we run the simulator with 64 threads. Our machine also has a GPU, an Nvidia Geforce GTX1080.

The simulation is a simple one: we illuminate a flat semi-infinite silicon sample with an electron beam of varying energy. We measure the throughput, in primary electrons per second, on both the CPU and GPU. This is shown in figure 3.4a. Figure 3.4b shows the speedup we get from the GPU.

A factor 2 speedup from the GPU may not sound very good, but this comparison is quite unfair. We are using two CPUs designed for computer clusters. Their retail price was more than €2100 each at the time of release in 2016. Our GPU is a gaming card for consumers which was also released in 2016, but it retailed for only €800. It would have been fairer to use either a high-end consumer CPU or a datacenter GPU for this comparison.




Figure 3.5: Energy spectrum of reflected electrons. The electron beam impinges perpendicularly on a flat, semi-infinite silicon sample with a 1 keV energy. Indicated are the contributions of backscattered electrons (BSE), true secondary electrons (SE), tertiary (and beyond) electrons (TE), and the sum (all). (a) zooms in on the lower energies, (b) shows the full spectrum.

3.3 Reflected electrons

We are now interested in an experiment in which a large and thick sample is illuminated with a 1 keV beam. We want to study the electrons emitted from such a sample on the “reflection” side. The term reflection here was borrowed from a common type of experiment known as reflection electron energy loss spectroscopy (REELS). The electrons do penetrate the sample in such an experiment; the term reflection reflects the fact that we are not looking at transmission through a thin film.

3.3.1 Energy spectrum

First, we investigate the energy spectrum of emitted electrons. We are not only interested in the total energy spectrum, but also in the fraction of backscattered primary electrons, secondary electrons generated by those primaries, and “tertiary electrons” which were generated by the secondaries. For simplicity, we will label all quaternary electrons and beyond as tertiary electrons. The energy spectrum is shown for a 1 keV electron beam on silicon in figure 3.5.

First, consider the backscattered electrons in figure 3.5b. Near 1000 eV, primary electrons enter the sample, lose a small amount of energy, and escape the

sample again. As plasmon creation is the most likely energy loss mechanism, the plasmon peak shows up at 17 eV energy loss. Instances of multiple plasmon creation show up as further peaks. They look like oscillations in figure 3.5b. At lower energies, the contribution of backscattered electrons diminishes. There is a small upturn again at very low energies, most likely due to the fact that the inelastic mean free path increases rapidly there.

The contribution of true secondary electrons is most significant at very low energies, figure 3.5a. The plasmon peak is clearly visible in their energy spectrum. Most secondary electrons are created from near the top of the valence band, because this is where the density of states is highest. With respect to this energy level, they gain the plasmon energy, which is 17 eV in silicon. This puts them 12 eV above the vacuum level, where 5 eV is lost due to the electron affinity and the band gap. Secondary electrons below the plasmon peak are somewhat overrepresented in figure 3.5a: they are not created as often, but they can escape more easily due to their longer mean free path. The majority of the energy spectrum at very low energies is made up of tertiary or higher order electrons. The plasmon peak is much less clearly visible in the total spectrum as a result. If one measured the energy spectrum of the electrons reaching the surface, which is not shown here, one would observe that the number of secondary and tertiary electrons keeps increasing towards lower energies. The characteristic “down-turn” near zero energy is caused by the work function surface barrier, which stops low-energy electrons escaping unless their orientation is just right.

3.3.2 Spatial distribution

We now turn our attention to the spatial distribution of emitted electrons. The primary electron lands at $(x, y) = (0, 0)$. We measure the (x, y) position where the electrons are emitted again. As in the previous section, we split up the signal into contributions from backscattered electrons (BSEs), true secondary electrons (SEs) and tertiary (or higher order) electrons (TE). Additionally, we split the contribution of SEs up into SE1 and SE2. When an SE was created by a primary electron with a downward direction (negative z component in its direction vector), we label it an SE1. If the primary electron had an upward direction, we label it SE2. This is to connect to literature, where SE1s are understood to be created by the primary beam going down and SE2s are created by backscattered electrons going up. The distributions can be seen in figure 3.6.

The highest density of BSEs is in the center, near the point where the primary electron struck. Their density decreases according to a power law before

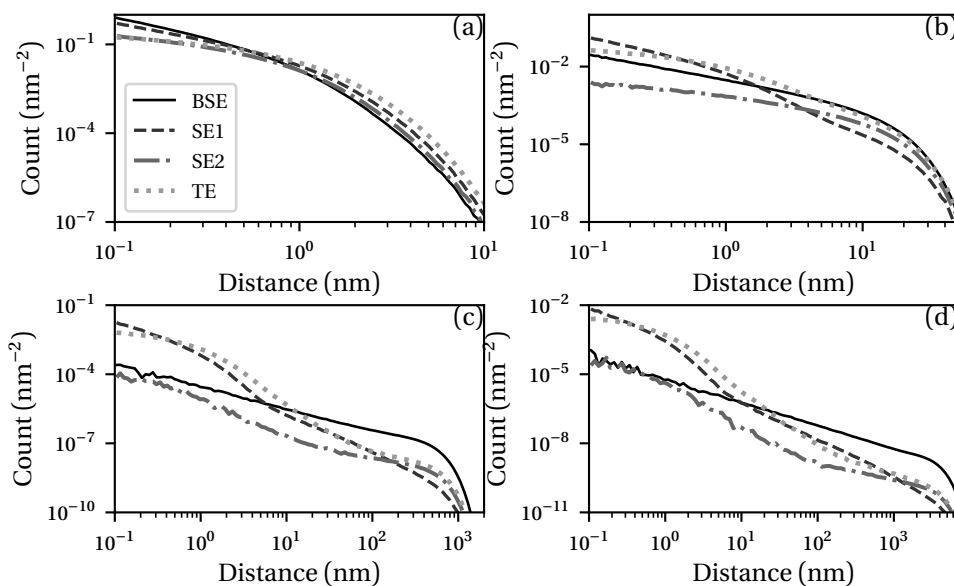


Figure 3.6: Spatial distribution of emitted electrons. The horizontal axes measure the distance from the primary electron's point of incidence. The vertical axes measure density of emitted electrons per unit area per primary electron. Curves are the contributions of backscattered electrons (BSE), true secondary electrons (SE), and tertiary (and beyond) electrons (TE). SEs are split up into SE1 and SE2, which were created by primaries with moving into the negative and positive z directions, respectively. The various figures are for different energies of the primary electron beam: (a) 100 eV; (b) 1 keV; (c) 10 keV; (d) 30 keV.

a sharp cutoff. Backscattered electrons may be emitted several μm from the primary electron's point of incidence if the landing energy is high enough. By contrast, SE1s are most prominent in the first few nm around the primary electron's point of incidence. After that, their density drops off steeply, but they do not disappear entirely until the backscatter range is reached. SE2s, which are created by BSEs, give a much less significant contribution than SE1s close to the primary impact point. Because they are created by BSEs, they have a much longer tail than SE1s and eventually start to dominate over SE1s. The density of TEs is roughly proportional to the density of SE1s plus that of SE2s, indicating that neither is particularly more effective at creating TEs.

3.4 Shape of the interaction volume

We are now interested in the shape of the interaction volume. Figure 3.7 shows maps of deposited energy as function of position. This figure counts “deposited energy” at the point where electrons lose energy. If a primary electron loses, for example, 30 eV creating a secondary electron, then the energy deposit is counted as 30 eV even though the secondary electron technically carries some of that energy away.

The primary electron beam enters the sample at the origin, $(r, z) = (0, 0)$, and travels in the $-z$ direction. It deposits a lot of energy near this point. Because the beam is oriented downwards, and scattering is mostly oriented in the forward direction, the beam only broadens a little bit in the initial scattering events. When the beam reaches a certain depth, it has lost this downward direction due to the many scattering events it has suffered. The depth at which this happens can be estimated by the elastic transport length of the primary electrons, figure 2.10b. This is a slight overestimation, because the transport length in figure 2.10b does not include the small-angle scattering from inelastic events, nor does it include the reduction in transport length due to the fact that the electrons lose energy. After this initial stage, the primary electrons move in a spherical shape around the depth given by the transport length. Electrons escaping through the surface do not come back. This leads to a decrease in deposited energy near the surface, and can be seen as the contour lines tend to be “pulled in” towards $r = 0$ near the surface.

A pear-like shape is often mentioned in literature [94–96]. Based on the explanation in the previous paragraph, this is true if we consider the electrons almost stopped when they have travelled their transport length. Indeed, the contour lines close to the origin in figures 3.7b–3.7d have pear-like shapes. If




Figure 3.7: Maps of deposited energy in a semi-infinite silicon sample. The electron beam enters the sample at the origin (0,0) pointing downwards. The horizontal axis shows the radial coordinate, the vertical axis shows the depth in the material. The colour scale is the amount of energy deposited per unit volume per primary electron, in eV nm^{-3} . The gray curves are contour lines of equal deposited energy. The black dashed lines indicate the depth above which 90% of all energy has been deposited. The various images are for different energies of the primary electron beam: (a) 100 eV; (b) 1 keV; (c) 10 keV; (d) 30 keV.

we also consider those few electrons that travel a much longer way, the interaction volume looks more like a hemisphere centred at a depth corresponding to the transport length, with a slight reduction near the surface due to escaping electrons. The amount of energy deposition in the hemisphere regime is orders of magnitude less than in the pear-shape regime. The pear-shape regime is therefore the most relevant in practical situations.

In the case of 100 eV landing energy, as in figure 3.7a, the interaction volume looks very different. We attribute this to the following. At 100 eV energy, the inelastic mean free path is much shorter than the electron–phonon mean free path (compare figures 2.8a and 2.13a). The primary electron changes direction in an inelastic event. This is usually a small amount, but if the primary electron only has 100 eV, it can be significant. A very rough estimate for the scattering angle of the primary electron can be found using equation (2.71)⁴. A 100 eV electron losing 17 eV (the plasmon energy in silicon) is deflected by approximately 34°. The primary electron’s trajectory is therefore significantly affected by inelastic scattering. The “center” of the hemisphere is therefore very close to the surface, much closer than the electron–phonon transport length of figure 2.13b suggests. The question remains why the contour lines in figure 3.7a “bend downwards” and run almost parallel to the $r = 0$ axis instead of intersecting it at right angles as they do for the other landing energies. It is likely that this is caused by primary electrons that did not scatter at all in the top nm. Hence, these energy loss events take place at exactly $r = 0$. This behaviour is very different from the figures at higher energy, where electrons have undergone many scattering events by the time the transport length is reached. Therefore, they are much more likely to have moved away from the central $r = 0$ axis.

3.5 Experimental yield comparison

Now that we have a fully functioning simulator, it is time to compare simulations to experimental results. One very easy way to do this is by comparing electron yields. The idea is simple: we take a very large, thick, flat sample and illuminate it with an electron beam. We let the beam’s landing energy vary. We then measure how many electrons are emitted from the sample in the direction of the beam.

⁴Because the target electron is not free, this is only a rough estimate. It would be better to determine the full angular distribution from the probability distributions for energy and momentum transfer, equation (2.34).

Two kinds of electrons are distinguished: backscattered electrons (BSEs), with energies greater than 50 eV, and secondary electrons (SEs), with energies less than 50 eV. Note that this definition of SE does not necessarily overlap with the notion of secondary electron in the sense of being an electron that comes from the material. It is quite possible that such a secondary electron has more than 50 eV energy. In practice, the distinction between SEs and BSEs makes sense because the dominant secondary electron excitation mechanism is plasmon decay, and secondary electrons with more than 50 eV are uncommon. Similarly, it is unlikely for a primary electron starting with many hundreds of eV to be emitted with less than 50 eV. By the time it has lost that much energy, it is most likely very deep below the surface and is unable to escape.

At this point, it should be noted that although there is a lot of experimental yield data available, not all of it is reliable. The condition of the surface, *e.g.* due to the presence of native oxides or hydrocarbons [97, 98] is an important factor. This is especially the case for SE yields: most SEs have energies less than 5 eV relative to the vacuum level. They originate from the top few nm of the sample and are easily affected by such thin layers. BSEs tend to come from deeper layers and have a higher probability of passing through a surface layer unaffected. Insulating samples are often heavily affected by charging. As a result, there is a significant amount of scatter in experimental data.

We have selected two studies which have taken good care to suppress experimental errors as well as possible. One is the study of Bronstein and Fraiman [99], which is frequently cited in literature. This reference is in Russian and is not readily available in some university libraries. Fortunately, the data is available in the database of Joy [53] and much of the relevant experimental information has appeared in various articles which have been translated to English [100–105]. The second source is the more recent study (published in two papers) of Walker *et al.* [106] and El Gomati *et al.* [107]. Our simulations are compared to the measurements in figure 3.8.

The general agreement between simulations and experiment is quite good. The most striking discrepancy exists for the SE yield of aluminium, figure 3.8c. Walker *et al.* find a much larger SE yield than Bronstein and Fraiman. According to Walker *et al.*, a native oxide on aluminium increases the SE yield, and it is particularly difficult to remove the native oxide from aluminium. They suggest that the differences between them and Bronstein and Fraiman may be due to instrumental variations, or, alternatively, that Bronstein and Fraiman may have been the only authors to achieve a truly clean surface. As we will show in section 6.8, the effect of an oxide layer on the SE yield can indeed be substantial.




Figure 3.8: Measured and simulated electron yields. Left column: SE yields, right column: BSE yields. Top row: silicon, middle row: aluminium, bottom row: copper. Curves are simulations, points are measured data. Closed symbols are from the study of Bronstein and Fraiman [99], open symbols are from Walker *et al.* [106] (for SE yield) and El Gomati *et al.* [107] (for BSE yield).




Figure 3.9: Modified (a) inelastic and (b) phonon mean free paths.

3.6 Sensitivity analysis

As discussed in chapter 2, there is considerable uncertainty in the physical models. This is especially the case at low electron energies. It therefore makes sense to study the sensitivity of the simulator to these uncertainties.

We perform the following study. We multiply the scattering mean free paths by an arbitrary factor, but only at low energies. The arbitrary factors we choose here are 2 and 1/2. For inelastic scattering we multiply the sub-100 eV energy region, leaving everything above this region untouched. The intention behind this is to capture the variation in mean free paths seen in figure 2.8a. We do not attempt to make this a smooth transition, hence there is an artificial step in the mean free path at 100 eV (see figure 3.9a). For elastic scattering, we have good faith in the high-energy Mott scattering model but not in the low-energy phonon model. Therefore, we multiply only the phonon model by an arbitrary factor. Because we already let the phonon model transition smoothly into the Mott model, the transition here is also smooth in the 100–200 eV region. See figure 3.9b.

3.6.1 Electron yields

First, we study the impact of this factor on electron yields. This can be seen in figure 3.10. When the inelastic mean free path is multiplied by 2 at low energies, the range that secondary electrons can travel is increased. This leads to a substantial increase of the SE yield (figure 3.10a), and the energy at which the yield




Figure 3.10: Top row: SE yields, using modified (a) inelastic and (b) phonon mean free paths as seen in figure 3.9. Bottom row: BSE yields, using modified (c) inelastic and (d) phonon mean free paths.

is maximal shifts up. The increase in SE yield is most significant at landing energies higher than 100 eV. This is because, at landing energies below 100 eV, mean free path of the primary electrons is also increased. The secondary electrons are therefore created deeper in the material, balancing the increased yield due to their larger range of travel. The fact that the energy of maximal yield shifts up can be explained as follows. The point of maximal yield is caused by a balance, where lower-energy primary electrons have less energy to create SEs in this surface layer and higher-energy primary electrons create too many SEs below the surface layer. The increased range of secondary electrons means that the depth from which they can escape increases. Hence, the maximum yield is found at a higher beam energy. When the inelastic mean free path is reduced by a factor 2, we see the opposite: both the SE yield and the energy of maximal SE yield decrease. Similar effects happen in the BSE yield (figure 3.10c), though not as severely. This is because of the electrons in the 50–100 eV energy range. These are often high-energy secondary electrons, but they are classified as BSEs due to their energy.

We now turn our attention to the SE yield in case the phonon mean free path

is modified, figure 3.10b. At high landing energies, increasing the phonon mean free path leads to a higher yield because secondary electrons can travel further. The behaviour at lower energies is more complex. The energy at which the yield is maximal shifts to a higher energy, but the maximum yield itself is decreased. It seems that the following effects are at play. Secondary electrons escape from a very thin layer near the surface. At very low energy, primary electrons have short mean free path lengths and they are deflected substantially in both inelastic and elastic events. They therefore create multiple secondary electrons within this surface layer. A higher beam energy increases the SE yield because there is more energy available to create secondary electrons. At higher beam energies, the mean free paths of primary electrons increase and the deflection in scattering events become less. The primary electrons penetrate through the surface layer and create most secondary electrons in a region from which they cannot escape. This leads to the maximum in SE yield. Increasing the electron–phonon mean free path does two things: it allows 0–200 eV primary electrons to penetrate deeper, and it allows the secondary electrons to escape more easily. Note that the net travel length for primary and secondary electrons does not change by the same amount because inelastic scattering is unaffected when we change the electron–phonon mean free path. Apparently, increasing the electron–phonon mean free path increases the travel range for secondary electrons less than it increases the travel range for primary electrons. The net result is that the yield is lowered for low landing energies.

3.6.2 Linescan

We now turn our attention to scanning electron microscopy. We are interested in the linescan signal of an isolated silicon line on a silicon substrate. The line is 40 nm high and 40 nm wide, we use an infinitely sharp electron beam and a beam energy of 500 eV. See figure 3.11. Each of the linescans follows a well-known pattern. On the flat parts of the sample, far away from the edge, the yield is almost independent of the beam position and the curve is flat. On top of the line, the yield is higher near the edge than far away from it. We call this increased yield near the edge the “edge blooming” effect. This is because secondary electrons may escape from the side of the structure. Immediately next to the line, when the beam lands on the substrate, the yield is considerably lower because escaping electrons are blocked by the line edge.

When increasing the inelastic mean free path at low energies (figure 3.11 a), the secondary electrons get a longer range of travel. This leads to a significantly




Figure 3.11: Linescan of an isolated 40 nm high and 40 nm wide silicon line on a silicon substrate. The feature is centered around $x = 0$. The electron beam is infinitely sharp and has a 500 eV landing energy. The intensity on the vertical axis is derived from the emitted secondary electrons. We modify the (a) inelastic and (b) phonon mean free paths as shown in figure 3.9.

increased yield, as seen before. As may be expected, the edge blooming effect simultaneously becomes wider. Similar effects happen when the phonon mean free path is increased (figure 3.11b), again, as expected.

We can quantify the width of the edge blooming effect as follows. We define a “base” grayscale value in the middle of the feature, *i.e.* at $x = 0$ nm in figure 3.11. The SE yield is maximal near the line edge, at $x \approx \pm 20$ nm in the figure. We find the x position where the grayscale value is halfway between these two values. We define the difference between this point and the x position of the maximum as the width of the edge blooming effect. The edge blooming widths are given in table 3.1. Reducing either the inelastic or phonon mean free path by a factor 2 leads to a 20% reduced edge blooming effect; increasing them by a factor 2 leads to a 25% increase.

3.6.3 Discussion

We can draw the following conclusions from this sensitivity analysis.

The secondary electron yield is very sensitive to low-energy scattering cross sections. A factor 2 in inelastic mean free path leads to a considerably different yield. A factor 2 in phonon mean free path has a much smaller impact, but it must be noted that the physical uncertainty in the phonon scattering mecha-

Table 3.1: Width of the edge blooming effect in figure 3.11.

Condition	Edge blooming (nm)
Baseline	2.73
Inelastic $\times 2$	3.39
Inelastic $/ 2$	2.23
Phonon $\times 2$	3.42
Phonon $/ 2$	2.21

nism is much larger than that of the dielectric function formalism. While it is likely that the error in the dielectric function mean free path is less than a factor 2, it may be larger than that for phonon scattering. Not only have we observed that the yield itself increases or decreases when the mean free paths are varied, the same is true for the energy where the yield is maximal. This is another common discrepancy between simulation and experiment. As shown in figure 3.8, the simulator tends to underestimate the energy where the SE yield is maximal. We can shift the energy of maximal yield up by increasing the range of travel for the secondary electrons, either from the phonon or the dielectric function side. This indicates that the range of SEs may be underestimated with our current models.

The backscattered electron yield is not very sensitive to the changes we have investigated. This is not surprising as our sensitivity analysis tackles only BSEs in the 50–100 eV range. One thing to note here is that the BSE yield has a “shoulder” at 200 eV which disappears when the phonon mean free path is increased (figure 3.10d). In section 6.3, we will see that this shoulder can also be removed by tweaking the 100–200 eV interpolation range between phonon and Mott scattering. Because this interpolation range was chosen arbitrarily, this shoulder is likely an artefact of our modelling rather than a physical effect.

The width of the edge blooming effect changes by about half a nm (for a 500 eV electron beam) when either mean free path is changed by a factor 2. The width of the edge blooming effect is a measure for the size of the interaction volume. Visual comparisons of our simulated images to real SEM images often suggest that the simulator underestimates the width of the edge blooming effect. This hints to the same conclusion as above: we may be underestimating the range of secondary electrons in our simulations. The result is that real SEM images may be slightly more “blurry” than what our simulator indicates.

Chapter 4

Surface effects

Monte Carlo simulations of SEM images ignore most surface effects, such as surface plasmons. Previous experiments have shown that surface plasmons play an important role in the emission of secondary electrons. We investigate the influence of incorporating surface plasmons into simulations of low-voltage CD-SEM. We use a modified inelastic scattering model, derived for infinite flat surfaces, and apply it to non-flat, but smooth, geometries. This simplification captures most qualitative effects, including both surface plasmons and a reduced interaction with bulk plasmons near interfaces. We find that the SE signal hardly changes when surface interactions are turned on for a perpendicularly incident beam. When the incident beam is perfectly parallel to a surface, the SE signal does significantly increase. However, the beam must be extremely close to the surface for this effect to be appreciable. A SEM is unable to produce a beam that is both narrow and parallel enough to be noticeably affected. The position of edges may appear shifted under specific circumstances. In realistic situations, it is unlikely to be a large effect.

*Parts of this chapter have been published as L. van Kessel, C.W. Hagen, and P. Kruit, "Surface effects in simulations of scanning electron microscopy images", Journal of Micro/Nanolithography, MEMS, and MOEMS **18**, 044002 (2019) [108].*

4.1 Introduction

Scanning electron microscopy is a standard tool for the inspection and metrology of semiconductor devices. It involves a focused electron beam scanning over a sample. The beam-sample interaction is a non-trivial process, producing secondary electrons (SEs) in a finite-sized interaction volume as well as

backscattered electrons (BSEs). The interpretation of a SEM image can be fairly straightforward if the features are large, but details of the beam–sample interaction become increasingly important as device features continue to shrink.

Semi-analytical models for the SEM signal have been proposed [109, 110], but the most rigorous method currently used is Monte Carlo simulations [9–17]. These simulations attempt to predict the SEM image, assuming the sample geometry and material composition are known exactly. Starting from physical principles, they make broadly similar assumptions. An electron is treated as a classical point particle, scattering through the volume of a material in discrete events. The electron is treated as if it is in free flight between such events. The scattering probability per unit distance travelled, $p(x)$, is given by an exponential distribution,

$$p(x) = \frac{1}{\lambda} e^{-x/\lambda}, \quad (4.1)$$

where λ is called the mean free path.

Two independent types of scattering are typically distinguished: elastic scattering, where an electron changes direction without losing energy; and inelastic scattering, where an electron loses energy and may excite a secondary electron. Each of these mechanisms has its own mean free path. In addition, as an electron reaches a material interface, it may be reflected or refracted due to the change of inner potential between the materials.

The reflection and refraction of electrons at a surface is only one of the many effects that play a role at an interface. For example, there may be oxidation layers, dangling bonds and contamination. Such highly sample-dependent effects are not considered in this work.

Instead, we correct the inelastic scattering mechanism for surface effects. An assumption that is often tacitly made, is that elastic and inelastic scattering always behave as if the electron is deep inside bulk material. For scattering mechanisms that can be well described by electrons scattering on isolated atoms, such as inner-shell excitation or elastic Mott scattering, this is a valid approximation. For events where the electron probes the solid-state bulk, such as electron–phonon scattering or plasmon excitation, we may expect scattering behaviour to be different near an interface.

Indeed, coincidence measurements [111–113] have provided evidence that surface plasmons may contribute significantly to SE emission when beam energies are on the order of 100 keV. Werner *et al.* [114] more recently performed a similar study at a beam energy (100 eV) close to low-voltage SEM. They find clear evidence for the contribution of surface plasmon decay to SE emission, of

similar magnitude to the volume plasmon. Neglecting surface plasmons in simulations of SEM images, as is typically done, may therefore not be acceptable.

The goal of the present work is to study the sensitivity of simulated SEM images to the inclusion of surface plasmons.

We follow a framework similar to previous simulations of SEM images, replacing the inelastic scattering mechanism by one where the material interface is explicitly taken into account. This dielectric formalism includes surface plasmons as well as the “Begrenzung” effect, the reduced coupling strength to the volume plasmon near an interface. The goal of the present work is to study the sensitivity of simulated SEM images to the inclusion of these effects.

We introduce surface plasmons to our simulator by replacing the inelastic scattering model. This is discussed in detail in section 4.2. Results are shown in section 4.3.

4.2 Surface formalism

Typical modern simulations use the dielectric function to describe the inelastic scattering of electrons in matter. The strength of this approach is that the dielectric function can be measured in the optical regime. Assuming that all of space is occupied by a material with dielectric function $\epsilon(q, \omega)$, and that the electron is a single point charge, Maxwell’s equations can be solved to find the induced electric field in the material. This electric field is oriented such that it slows down the electron, which is interpreted as inelastic scattering.

The result is an inelastic mean free path given by equation (2.34),

$$\lambda^{-1} = \frac{1}{\pi a_0 E} \int d\omega \int \frac{dq}{q} \operatorname{Im} \left[\frac{-1}{\epsilon(q, \omega)} \right]. \quad (4.2)$$

Here, ω represents the energy loss, q is the momentum transferred, E is the primary electron’s kinetic energy before the collision and a_0 is the Bohr radius. Integration is over the kinematically allowed range. Probability densities for energy and momentum loss are given by $\lambda \partial \lambda^{-1} / \partial \omega$ and $\lambda \partial \lambda^{-1} / \partial q$, respectively.

In the general case, when not all of space is occupied by the same material, a similar analysis can be performed. However, the technicalities of solving the electric field are much more complicated. This has been done analytically for simple geometries [115], such as infinite flat surfaces, wedges, or spheres.

For structures relevant to lithography, finding analytical solutions is impractical or impossible. A physically accurate alternative would involve numerically




Figure 4.1: Geometry used by Salvat-Pujol and Werner [116]. Shaded areas represent the two materials, represented by dielectric functions $\epsilon_a(\mathbf{q}, \omega)$ and $\epsilon_b(\mathbf{q}, \omega)$. They are separated by the $z = 0$ plane. The electron moves with velocity \mathbf{v} and passes through the origin at time $t = 0$.

solving the induced electric field for every electron position, direction and energy of interest. This parameter space is prohibitively large.

We make the following simplification. We use the analytically known solution for infinite flat surfaces, and apply it to arbitrary geometries assuming that the radius of curvature in the geometry is sufficiently large everywhere. We will investigate the meaning of “sufficiently large” in more detail later.

We use the results from Salvat-Pujol and Werner [116]. A full derivation is given in that reference, we only repeat the main results. The situation is illustrated in figure 4.1. Materials with dielectric functions $\epsilon_a(\mathbf{q}, \omega)$ and $\epsilon_b(\mathbf{q}, \omega)$ fill the spaces $z > 0$ and $z < 0$, respectively. The electron moves with velocity \mathbf{v} . The electron’s z coordinate is denoted as d , positive for region a and negative for region b . At time $t = 0$, the electron passes through the origin. The electron’s position is given by $\mathbf{x} = \mathbf{v}t$.

Salvat-Pujol and Werner solve this situation for a non-relativistic electron by means of the image charge method. The assumption that the electron is not relativistic makes the problem formally electrostatic. Salvat-Pujol and Werner remark that the usual boundary conditions for the electric and displacement fields at $z = 0$ are not restrictive enough to uniquely establish the image charges. Various configurations for the image charges have been chosen in previous literature. Salvat-Pujol and Werner parametrise the choice for image charges by the

tuple of numbers (p_1, p_2, p_3) , which can be set to $(0, 1, 0)$ or $(1, 1, 1)$ to reproduce earlier literature results.

The result is an inverse mean free path consisting of a bulk and a surface term: $\lambda^{-1} = \lambda_B^{-1} + \lambda_S^{-1}$. The surface term captures both the surface plasmon coupling as well as the ‘‘Begrenzung’’ of the volume plasmon. The bulk term, λ_B^{-1} , is similar to equation (4.2), where ϵ_a or ϵ_b is selected depending on the sign of d . The surface interaction is given by

$$\begin{aligned} \lambda_S^{-1} = & -\frac{\hbar}{m_e a_0 \pi^2} \frac{|v_z|}{v} \int_0^{E/\hbar} d\omega \int_{q_-}^{q_+} dq \int_0^\pi d\theta \sin\theta \int_0^{2\pi} d\phi \\ & \frac{1}{q_\parallel^2 v_z^2 + (\omega - \mathbf{q}_\parallel \cdot \mathbf{v}_\parallel)^2} \\ & \times \left[\Theta(-t) e^{-q_\parallel |d|} + \Theta(t) \left(2 \cos \left\{ (\omega - \mathbf{q}_\parallel \cdot \mathbf{v}_\parallel) t \right\} - e^{-q_\parallel |d|} \right) \right] \\ & \times \text{Im} \left[e^{iq_z d} f(\mathbf{q}_\parallel, \omega) \left\{ \frac{\Theta(d)}{\epsilon_a(\mathbf{q}, \omega)} - \frac{\Theta(-d)}{\epsilon_b(\mathbf{q}, \omega)} \right\} \right]. \end{aligned} \quad (4.3)$$

We use the notation that any vector \mathbf{a} has components $\mathbf{a} = (a_x, a_y, a_z)$, the shorthand $a = |\mathbf{a}|$, and $\mathbf{a}_\parallel = (a_x, a_y)$ is the two-dimensional projection oriented parallel to the interface. Furthermore, $\mathbf{q} = q(\cos\phi \sin\theta, \sin\phi \sin\theta, \cos\theta)$, $\Theta(x)$ is the Heaviside step function, and

$$\frac{\hbar q_\pm}{\sqrt{2m_e}} = \sqrt{E} \pm \sqrt{E - \omega}, \quad (4.4)$$

$$k_z = \frac{\omega - \mathbf{q}_\parallel \cdot \mathbf{v}_\parallel}{v_z}, \quad (4.5)$$

$$f(\mathbf{q}_\parallel, \omega) = \frac{\frac{p_2}{\epsilon_b(\mathbf{q}_\parallel, k_z, \omega)} - \frac{1}{\epsilon_a(\mathbf{q}_\parallel, k_z, \omega)} + \frac{p_3}{\epsilon_b(\mathbf{q}_\parallel, -k_z, \omega)} - \frac{p_1}{\epsilon_a(\mathbf{q}_\parallel, -k_z, \omega)}}{\int_{-\infty}^{\infty} d\kappa \frac{1}{q_\parallel^2 + \kappa^2} \left[\frac{1}{\epsilon_a(\mathbf{q}_\parallel, \kappa, \omega)} + \frac{1}{\epsilon_b(\mathbf{q}_\parallel, \kappa, \omega)} \right]}. \quad (4.6)$$

As before, the probability distributions of losing energy ω or momentum \mathbf{q} are given by $\lambda \partial \lambda^{-1} / \partial \omega$ and $\lambda \partial \lambda^{-1} / \partial \mathbf{q}$.

We emphasise that λ_S^{-1} cannot be seen separately from λ_B^{-1} . λ_S^{-1} can be negative, representing a reduction in the bulk interaction. The sum $\lambda^{-1} = \lambda_B^{-1} + \lambda_S^{-1}$, however, must be positive.

We also note that λ_S^{-1} changes as the electron travels. λ^{-1} represents the scattering cross section in the electron’s immediate environment, which is usually constant between successive events. This is not true in the present case, so

λ can no longer be directly interpreted as the mean distance travelled between scattering events.

4.2.1 Features of the surface correction

A thorough discussion of the physical features contained in the surface formalism is already given by Salvat-Pujol and Werner [116]. For the reader's convenience, we repeat some of the main conclusions. We use the same dielectric function,

$$\frac{1}{\epsilon(q, \omega)} = 1 - \frac{\Omega_p^2}{\hbar^2 Z} \sum_{j=1}^n \frac{f_j}{\omega_j^2 + \hbar^2 q^4 / 4m_e^2 - \omega^2 - i\gamma_j \omega}, \quad (4.7)$$

with, for aluminium, $Z = 13$, $\Omega_p = 32.84$ eV, $n = 1$, $\hbar\omega_1 = 15.01$ eV, $f_1 = 3$, $\hbar\gamma_1 = 0.5$ eV. This facilitates direct comparison to results shown by Salvat-Pujol and Werner.

Later on in this work, we will perform Monte Carlo simulations for silicon. We use a more realistic dielectric function for that. We have obtained the energy-loss function $\text{Im}[1/\epsilon(q, \omega)]$ from measured data at $q = 0$ [29], which is extended into the $q > 0$ regime by means of the full Penn algorithm [42]. Our numerical implementation of the Penn algorithm follows the description of Shinotsuka *et al.* [26]. The real part of $1/\epsilon(q, \omega)$ is then obtained by means of the Kramers-Kronig relation

$$\text{Re} \left[\frac{1}{\epsilon(q, \omega)} \right] = 1 + \frac{1}{\pi} \mathcal{P} \int_{-\infty}^{\infty} \frac{d\xi}{\xi - \omega} \text{Im} \left[\frac{1}{\epsilon(q, \omega)} \right], \quad (4.8)$$

where \mathcal{P} stands for the principal value of the integral. We numerically evaluate this equation using the double Fourier transform method [117].

Let us now consider an electron landing perpendicularly on a flat vacuum–aluminium interface. Figure 4.2 shows the differential inverse mean free path (DIMFP) with respect to energy, $\partial\lambda^{-1}/\partial\omega$, for a 100 eV electron for several distances d to the interface. In figure 4.2a, the electron is on the vacuum side; in figure 4.2b, it is on the aluminium side.

When the electron is far on the vacuum side, it can't lose any energy. However, as it approaches the interface, a peak appears near 11 eV, aluminium's surface plasmon energy. The surface plasmon coupling increases in strength as the electron closes in on the surface.

As soon as the electron moves to the aluminium side, a second peak appears beside the surface plasmon interaction. This second peak represents the volume plasmon of aluminium. As the electron moves deeper into the material,




Figure 4.2: Differential inverse mean free path (DIMFP) for a 100 eV electron, as a function of energy loss ω , for several distances d to a vacuum–aluminium interface. This represents the probability for an electron to lose a certain amount of energy, per unit distance travelled. The electron travels from the vacuum towards the aluminium, along the surface normal. In (a), $d > 0$, the electron is on the vacuum side; in (b), it is on the aluminium side. The $d = -0.5 \text{ \AA}$ curve in (b) is hidden behind the curve for $d = -1 \text{ \AA}$.




Figure 4.3: Same as figure 4.2, for a 500 eV electron.

the surface plasmon coupling decreases in strength while the volume plasmon becomes stronger. This reduction of the volume plasmon coupling near the interface is known as the “Begrenzung” effect. At a depth of 0.5 nm, the DIMFP is very similar to the bulk DIMFP.

Similar effects can be seen for a 500 eV electron (figure 4.3). For the purposes of this work, we want to note two things. First, the “surface layer” is somewhat thicker. For a 100 eV electron, the surface interaction becomes negligible at less than 5 Å from the interface. This interaction extends further for a 500 eV electron. In general, the surface layer is thicker for faster electrons. Second, the 500 eV electron has a lower interaction probability per unit distance travelled. This is a familiar effect also observed in bulk mean free paths.

We emphasise the extremely short range in which the surface correction is relevant. Even for 500 eV electrons, the surface layer is less than a nm thick. Figures 4.2 and 4.3 show data for distances of 0.5 Å from the surface, much less than the distance between two atoms¹. The dielectric function is a continuum approximation to the microscopic response of the material, and whether it is a valid means to treat surface interactions at such small scales (low electron energies) is open for debate.

4.2.2 Justification of the infinite flat plane approximation

We have assumed that a curved geometry can be approximated as flat, if the curvature is smooth enough. This assumption needs to be justified, which we will do in two ways.

First, we show the lateral extent of the induced surface charge. In the work of Salvat-Pujol and Werner, there are three types of contributions to the electric field: the electron itself, its image charges in the bulk, and the induced surface charge. Clearly, if the geometry curves significantly on a scale where the surface charge is significant, the boundary conditions at the interface are not satisfied and the flat-plane approximation is wrong. Conversely, we may hope that if the interface curves far away from any significant surface charge, the boundary conditions at the interface are still approximately satisfied.

One might naively expect that the lateral extent of the surface charge is similar to the electron’s distance from the surface. Since the surface interaction vanishes for distances less than a nm, one may expect the minimum radius of curvature to be similar. However, it is good to verify this.

Figure 4.4 shows the surface charge for 300 eV and 500 eV electrons, each

¹Salvat-Pujol and Werner [116] dare to go even lower.




Figure 4.4: Induced surface charge on a flat aluminium surface, for (a) a 300 eV and (b) a 500 eV electron. The electron is 5 \AA outside the aluminium, and travels towards the surface along the normal. The color scale has units of elementary charges per nm^2 .

5 \AA outside aluminium, moving directly towards the aluminium. This distance was chosen as it is the largest distance where the surface interaction is still noticeable. The surface charge can be seen to decay to almost zero within a few nm. When applying the surface formalism to curved geometries, we will use electrons of no more than 500 eV and a radius of curvature of at least 5 nm.

More justification comes from a numerical comparison to an analytical result for a curved geometry. García de Abajo [115] gives analytical expressions for several geometries. We study the case of an electron passing outside a spherical nanoparticle.

It is well-established that surface plasmons on spherical nanoparticles can behave qualitatively differently than on a flat plane. For a flat metallic plane in vacuum, the surface plasmon energy $\omega_s = \omega_p / \sqrt{2}$, where ω_p is the bulk plasmon energy. A nanosphere, however, can support a large spectrum of modes between $\omega_p / \sqrt{3}$ and $\omega_p / \sqrt{2}$. High-energy (of order 100 keV) electrons dominantly excite the low-order modes.

The situation is as follows. An electron is infinitely far away from a sphere with radius a . It moves past the sphere, with closest radius of passing $b > a$, and goes on to infinity. García de Abajo [115] then gives an expression for the “loss probability”, given per unit of transferred energy. Its integral is the total probability that the electron loses energy.

With the flat-plane approximation, we can replicate this situation with our

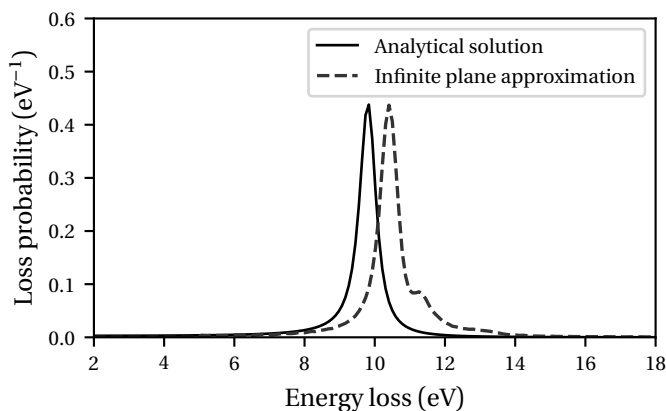


Figure 4.5: Loss probability distributions for a 500 eV electron passing at $b = 10.5$ nm from a $a = 10$ nm aluminium sphere. This figure compares the analytical solution (solid curve) to the result from the flat-plane approximation (dashed). Note the good agreement in absolute value.

Monte Carlo simulator. We start the electron sufficiently far away from the sphere, tracking its energy as it passes the sphere, recording its final energy when it has moved sufficiently far away. This gives a distribution that is directly comparable, both in the distribution of energy lost and in terms of absolute value, to the analytical result.

The analytical result predicts that the low-order plasmon modes dominate when the electron energy is large and the electron passes far away from the sphere. García de Abajo confirms this for the classic case of a 200 keV electron, and the parameters $(a, b) = (5, 6)$ nm or $(10, 12)$ nm. However, as mentioned, the much lower-energy electrons considered in this work barely interact with surface plasmons at such distances from the surface. They need to pass much closer to the surface. In the analytical result, the low-order modes are then suppressed, and higher-order modes become more important.

We show numerical results for the energy-loss distribution in figure 4.5. The situation is a 500 eV electron, passing a sphere with $a = 10$ nm at a distance $b = 10.5$ nm. The low-order modes in the analytical result are suppressed, and barely contribute to the peak in figure 4.5. Instead, the peak is caused by a combination of several higher-order modes. The amount of energy loss predicted by the analytical formula is slightly different from the flat-plane approximation. However, the agreement in absolute value is very good.

In conclusion, the amount of energy loss is slightly overestimated by the flat-

plane approximation. However, the number of events is captured very well. The latter property is the most important for the purpose of SEM images.

4.2.3 Monte Carlo implementation

The Monte Carlo implementation is not as straightforward as it is for bulk inelastic scattering. For bulk inelastic scattering from equation (4.2), one needs to store two- and three-dimensional tables. The two-dimensional table contains a probability distribution of energy loss, for each electron energy of interest. The three-dimensional table contains, for each electron energy and energy lost, a probability distribution of the momentum loss.

Surface inelastic scattering also depends on the electron's distance and angle to the surface, adding two dimensions to the tables mentioned above. In addition, we now need to sample two degrees of freedom for the momentum transfer. The bottom line is a need for 4, 5 and 6-dimensional tables: too large to fit in our computer's memory.

Therefore, we make the simplification that an inelastic event only causes the electron to slow down. We then only need a 4-dimensional table to sample the energy loss given the original energy, distance and angle to the surface.

We assume the following behaviour for the creation of SEs. When an inelastic event takes place, a secondary electron is created moving in a uniformly-distributed random direction. If the primary electron lost energy ω , the SE's energy is $E_0 + \omega$. E_0 is determined from the probability distribution $P(E_0, \omega) \propto \sqrt{E_0(E_0 + \omega)}$, where E_0 is between 0 and the Fermi energy. This probability distribution is the product of the SE's densities of states before and after the excitation. If the primary electron is in vacuum, the SE is started on the mirrored side of the boundary, inside the material. We acknowledge that this model is very simplistic and that there is room for improvement.

What remains to be discussed is when an inelastic event takes place. Remember that the "mean free path" changes as a function of the electron's position. A random free path length sampled according to equation (4.1) is therefore only correct near the electron's starting position. This can be solved by introducing a "maximum step length" between successive scattering events. This maximum step length should be small enough, such that the mean free path barely changes over this distance. We have used 0.1 Å here, but a larger value is possible without significantly influencing the results.

If the free path sampled according to equation (4.1) is longer than this maximum, we travel precisely the maximum and sample a new free path length. This

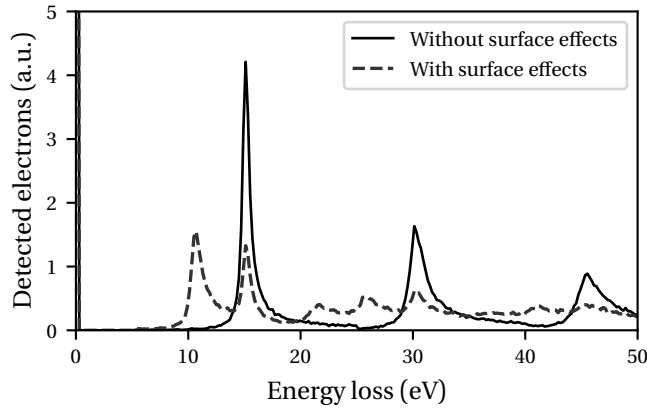


Figure 4.6: Energy loss spectrum, for a 300 eV beam on aluminium. Simulations with (dashed curves) and without (solid) surface effects enabled are compared.

mechanism prevents steps between the “physical” elastic or inelastic events from getting too large, and guarantees that the inelastic mean free path is suitably updated along the electron trajectory. Salvat-Pujol and Werner [116] do something similar, except their “null” event takes place stochastically with a mean free path λ_{\min} . This should lead to similar behaviour in the limit that λ_{\min} is much smaller than the elastic or inelastic mean free paths².

4.3 Results

4.3.1 Backscattered and secondary electron energy spectra

We present the simulated reflection electron energy loss spectrum (REELS) in figure 4.6. This was made for a 300 eV beam impinging perpendicularly on aluminium. Two simulations are shown: one with and one without surface effects enabled.

The energy spectrum without surface effects enabled is easily interpreted. The sharp peak at zero loss represents electrons reflected without losing energy. The strong peak at 15 eV represents electrons that have lost energy to a volume plasmon. Further peaks represent the excitation of multiple volume plasmons.

Enabling surface effects has a clear impact on the energy spectrum. The strong volume plasmon peak at 15 eV is reduced in absolute magnitude, and is

²There appears to be a small error in [116, step 3a]: λ_{\min}^{-1} should be replaced by $\lambda_{\min}^{-1} + \lambda_e^{-1} + \lambda_i^{-1}$. This is likely to be a typo, since the reference they provide for the algorithm is correct.




Figure 4.7: Secondary electron energy spectrum, for the same simulation as figure 4.6. Simulations with (dashed curves) and without (solid) surface effects enabled are compared. The material used is aluminium.

joined by the surface plasmon at approximately 11 eV. Subsequent peaks are caused by the excitation of multiple surface and/or volume plasmons. Clearly, if the backscattered energy spectrum is of interest, it is important that surface plasmons are taken into account.

The corresponding energy spectrum for secondary electrons is presented in figure 4.7. The distinctive peaks in the energy loss spectrum of figure 4.6 are not visible. No additional “surface plasmon” peak is visible when surface effects are enabled in the simulation. This is because of our assumptions in the SE generation mechanism: SEs generated after surface plasmon decay are created inside the material and have to overcome the work function barrier to escape the material. The probabilistic transmission model used at material boundaries (see section 2.3) effectively smooths out the features in the SE energy spectrum inside the material.

4.3.2 Yields

When taking SEM images, one is not typically interested in the detailed energy spectrum. Instead, detected electrons are counted, with only a very crude energy filter to distinguish between “secondary” (< 50 eV) and “backscattered” (> 50 eV) electrons. These yields—the average number of secondary or backscattered electrons per incident electron—are shown in figure 4.8 for silicon, as a function of beam energy.




Figure 4.8: Simulated (a) secondary and (b) backscattered yields, for simulations with (dashed curves) and without (solid) surface effects enabled. The material used is silicon.

These yields are similar to typical measured data [53] (not shown). Measured BSE yields saturate around 0.2. Measured SE yields tend to reach a maximum of 1.2 at a beam energy of approximately 300 eV, but there is a large (factor 2) spread between various experiments. SE yields are influenced by, amongst others, surface oxidation, contamination, dangling bonds, and surface roughness. None of these effects are present in these simulations. Interestingly, we find that the SE yield is barely changed when our surface model is enabled.

Section 4.2.3 mentions several simplifications that may influence the SE yield. These include the energy and direction of SEs, as well as their starting position. It is possible that a different implementation gives a slightly different SE yield when surface effects are enabled. However, it is unlikely to change the general picture much.

4.3.3 SEM image

We will now investigate the effect of introducing surface effects on a SEM image of a single silicon step (see figure 4.9). This step has circular rounded top and bottom corners to ensure smoothness. We vary the sidewall angle and electron beam energy. The electron beam is infinitely sharp, and we show only the SE signal.

We want to compare simulations with surface effects enabled and with surface effects disabled. We apply a vertical offset and linear scaling to the absolute




Figure 4.9: Step geometry. The shaded area represents the silicon. The top and bottom corners are circular, and the wall may be angled.

SE yield, such that the signals overlap away from the edge. This facilitates comparison of the signals near the edge.

We show the case of a vertical sidewall in figure 4.10. The step height is 20 nm and the corner radius is 5 nm, the beam energies are 300 eV and 500 eV. All SEM traces show the well-known enhancement of signal near the edge, which we will call the “edge blooming” effect in this work. When surface effects are enabled in the simulation, the SEM image taken at 300 eV shows a very sharp additional spike near the edge. This is a direct consequence of a somewhat pathological set of assumptions and parameters: the infinitely sharp electron beam travels at extremely close range ($x = 0.05$ nm) to the edge, and is oriented perfectly parallel to the edge. The electrons travel on the vacuum side and are able to efficiently excite surface plasmons, and the corresponding secondary electrons, without being deflected.

No spike can be seen when the beam energy is 500 eV, though the SEM signal is slightly enhanced at $x = 0.05$ nm when surface plasmons are enabled. The fact that the spike is not present in the 500 eV linescan is due to the larger mean free path for these higher-energy electrons.

When the step height is increased to 40 nm (figure 4.11), a spike also appears in the 500 eV SEM signal. The spike for a 300 eV beam at a 40 nm step is even higher.

It is impossible to create the exact circumstances of figures 4.10 and 4.11 in a regular CD-SEM. Still, if it were possible to perform this experiment, we expect such spikes to appear. The magnitude of the spikes, however, is likely to be overestimated by our model. Remember that a “surface event” does not deflect the primary electron in this model. The slight deflection that the primary electron suffers in reality would rapidly steer it out of the very narrow surface

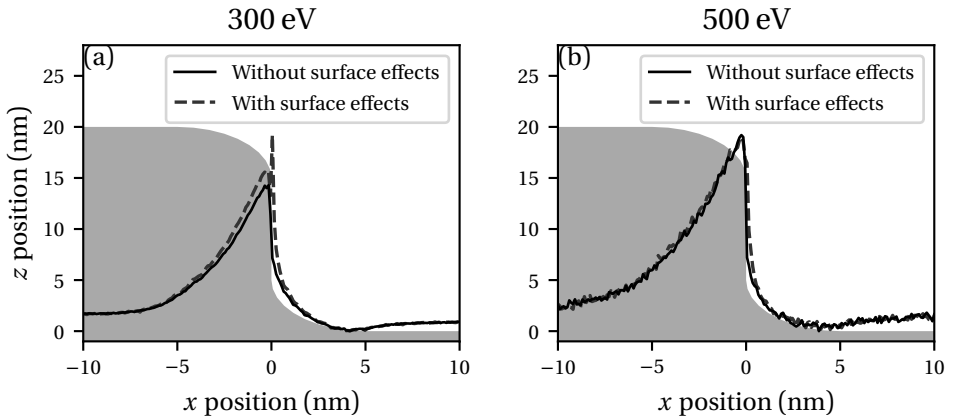


Figure 4.10: SEM signals for a geometry with $h = 20$ nm, $r = 5$ nm, $\alpha = 0^\circ$. Simulations with (dashed curves) and without (solid) surface effects enabled are compared. (a): beam energy 300 eV; (b): beam energy 500 eV. The shaded area represents the silicon feature. The feature appears skewed because the axes are not equally scaled.

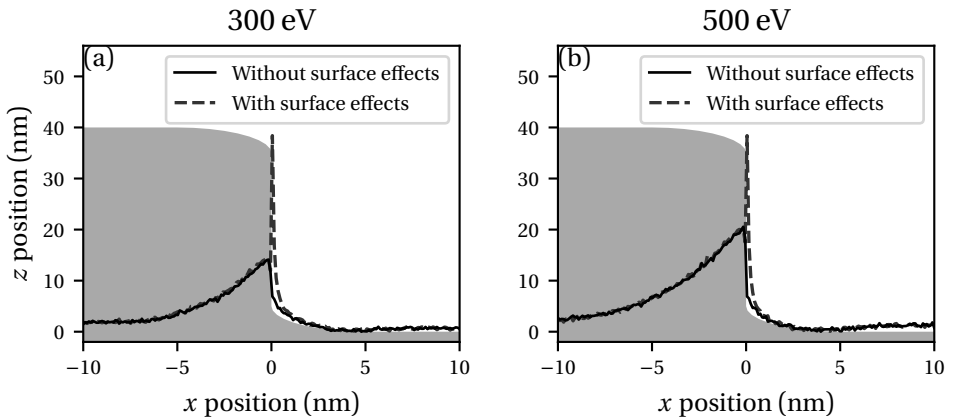


Figure 4.11: Same as figure 4.10, for $h = 40$ nm.

layer.

A 5° sidewall angle (figure 4.12) eliminates the observed spike. Instead, the SEM signal from the sidewall is somewhat enhanced when the electron beam is set to 300 eV. Surface effects make almost no difference under a 500 eV beam.

A feature with 1° sidewall angle (figure 4.13) holds few surprises given the other results. The signal from the sidewall-angle is enhanced, to a degree more extreme than seen in figure 4.12 but less than the spikes in figures 4.10 and 4.11. The simulation with surface effects enabled predicts that the signal at the sidewall exceeds the familiar edge blooming effect if the beam energy is low. As before, we expect the fact that the sidewall signal is enhanced to be true. The magnitude of the effect may be overestimated due to the contribution of multiple scattering events as well as the assumption that every surface plasmon decays to an electron.

The result in figure 4.13 is of interest to metrology. When surface effects are taken into account, the signal from the sidewall exceeds the edge blooming effect. This introduces a bias into the measured position of the edge.

There are multiple ways to measure the edge position in a SEM signal. A typical way is to take the position where the signal intensity is at 60% between the minimum and maximum. If the edge position is measured in this way, figures 4.10, 4.11 and 4.12 correspond to a bias of 0.1 nm or less. In figure 4.13, however, the simulations with surface effects included position the edge 0.6 nm further to the right than the corresponding simulations without surface effects included.

4.4 Discussion

At first glance, results obtained here seem to contradict experimental evidence mentioned in the introduction [111–114]. These experiments indicate that surface plasmon decay is an important contributor to SE emission.

The present simulations do not contradict these experiments. A SEM image counts the number of emitted SEs. As a consequence, information about the origin of SEs is lost.

Salvat-Pujol [118] performed coincidence experiments on aluminium and silicon. These experiments correlate the secondary electron's energy to the energy lost by the primary electron. They compared these experiments to their Monte Carlo model, which is very similar to ours. They show that reasonable agreement between simulation and experiment can only be obtained if surface plasmon decay is included in the model. In other words, the model captures the physics of SE generation by surface plasmon decay. In some situations, one




Figure 4.12: Same as figure 4.10, for $\alpha = 5^\circ$, $h = 40$ nm.




Figure 4.13: Same as figure 4.10, for $\alpha = 1^\circ$, $h = 40$ nm.

has no choice but to include surface plasmons in the model. For the purpose of SEM images however, they may be ignored.

4.5 Conclusions

The results shown above can be summarised as follows. Inclusion of surface effects in a Monte Carlo simulation leads to

- virtually no change in BSE yield,
- the SE yield changes only slightly,
- a changed BSE energy spectrum,
- a changed, but still mostly featureless, SE energy spectrum,
- a relative increase of the SE signal near a feature side-wall, which is largest for low landing energies.

Amongst others, we have assumed that incident electrons, upon coupling to a surface plasmon while in vacuum, do not change direction. We have also assumed that every surface plasmon decays to an electron-hole pair. These two assumptions lead to an overestimation of the amount by which the SE signal is increased near a side-wall.

In some of the situations shown, the enhancement of the signal from the wall is so strong that the edge blooming effect is surpassed. It is unclear whether this is realistic. It would be difficult to observe this effect directly: the “wall feature” in figure 4.13 is likely to appear indistinguishable from the edge blooming effect.

If the enhancement of the signal from the sidewall is indeed as big as these simulations suggest, this can lead to a different assignment of the edge position when the SEM image is interpreted. Effectively, this position may shift from the “top” of the sidewall to the “bottom” side.

Surface effects have a much less significant influence on the SEM signal near the rounded top and bottom corners of our feature. We acknowledge that our corner radius is much larger than typical. However, under the present physical assumptions, decreasing this radius does not lead to more interesting features in the SEM signal.

It is possible that more interesting features come to light when the flat-plane assumption is lifted and Maxwell’s equations are solved for the actual geometry

of interest. However, given that our present model captures most of the interesting features, we do not expect the qualitative conclusions to change much. The only interesting feature we observe, for the purposes of CD-SEM metrology, is an enhancement in signal from the sidewall angle. A more thorough treatment may, therefore, not be worth the effort.

Chapter 5

The influence of sidewall roughness on observed line-edge roughness in SEM images

Line-edge roughness (LER) is often measured from top-down critical dimension scanning electron microscope (CD-SEM) images. The true three-dimensional roughness profile of the sidewall is typically ignored in such analyses. We study the response of a CD-SEM to sidewall roughness (SWR) by simulation. We generate random rough lines and spaces, where the SWR is modelled by a known power spectral density. We then obtain corresponding CD-SEM images using a Monte Carlo electron scattering simulator. We find the measured LER from these images, and compare it to the known input roughness. For isolated lines, the SEM measures the outermost extrusion of the rough sidewall. The result is that the measured LER is up to a factor 2 less than the true on-wafer roughness. The effect can be modelled by making a top-down projection of the rough edge. Our model for isolated lines works fairly well for a dense grating of lines and spaces, as long as the trench width exceeds the line height. In order to obtain and compare accurate LER values, the projection effect of SWR needs to be taken into account.

Parts of this chapter have been published as L. van Kessel, T. Huisman, and C.W. Hagen, "Understanding the influence of three-dimensional sidewall roughness on observed line-edge roughness in scanning electron microscopy images", Journal of Micro/Nanolithography, MEMS, and MOEMS 19, 034002 (2020) [119]. It received the Karel Urbánek best student paper award when presented at the SPIE Advanced Lithography conference 2020.

5.1 Introduction

As lithography techniques move towards smaller features, line-edge roughness (LER) becomes an increasingly important parameter. Consequently, accurate metrology of LER has generated plenty of interest in recent years [120–122].

LER is usually measured from top-down critical dimension scanning electron microscope (CD-SEM) images. As CD-SEM is a two-dimensional imaging technique, the profile along the vertical direction is not easily accessible. Most studies of LER metrology by SEM therefore ignore the 3D roughness profile along the sidewall's vertical direction. In this work, we will use the term LER when referring to the roughness of the measured line profile in a CD-SEM image. Sidewall roughness (SWR) will refer to the physical roughness of the on-wafer profile.

SWR can be measured by, for example, tilted or cross-section SEM, or atomic force microscopy (AFM). Thick resist films exhibit near-isotropic SWR after development [123, 124]. Strong anisotropy is typically observed after etch [123, 125, 126]. Thin resists used for recent technology nodes show strong anisotropy with striations in the vertical direction [127, 128].

The influence of SWR on CD-SEM images has been experimentally investigated by Foucher *et al.* [124, 129] and Fouchier *et al.* [130, 131]. These experimental studies show that 3σ LER measured from CD-SEM can be significantly less—by more than a factor 2—compared to the 3σ SWR measured by AFM. Fouchier *et al.* [130, 131] conclude that the LER measured by CD-SEM matches the SWR measured by AFM very well, if the roughest upper layers of the feature are discarded from the AFM data. However, it is likely that the SEM is most sensitive to the upper layers, in which case figure 8 from [130] indicates the difference could be a factor 2.

Simulation studies have also been performed [91, 132, 133]. These studies start with a known geometry and obtain corresponding CD-SEM images using full electron scattering simulations. These simulation studies all confirm that the measured 3σ LER can be up to a factor 2 less than the true 3σ SWR. However, none of them attempt to explain where this bias comes from, or how the measured power spectral density (PSD) corresponds to the true physical PSD.

In this study, we model the most dominant effect for the bias between LER and SWR, and show how the observed PSD is related to the physical PSD. We do this by simulation. We generate line-space patterns with random rough sidewalls using a model PSD. We then obtain corresponding top-down SEM images using a Monte Carlo electron scattering simulator. We find the measured LER




Figure 5.1: Sketch of the line profile studied in this work. We start with a trapezoidal structure. The sidewalls are randomly displaced by a roughness model, while the top and substrate remain flat. Also shown is our choice of coordinate system and definition of sidewall angle α . In the main text, we only consider lines with $\alpha = 90^\circ$. Different sidewall angles are used in section 5.6.1.

by extracting the contours from these images, computing the measured PSD and subtracting the white noise floor. This unbiased measured roughness can then be compared to the input roughness.

The method of generating random features, simulating a SEM image, and extracting contours, is described in section 5.2. We first show results for the simple case of no roughness in the vertical direction in section 5.3. In section 5.4, we move to isolated lines with 3D roughness. Dense line-space patterns are studied in section 5.5.

5.2 Method

We study a pattern of lines and spaces in resist. We start with perfectly trapezoidal lines. The sidewalls of these lines are then modulated by random roughness according to a model PSD (section 5.2.1). To guarantee that the geometry never self-intersects, there is no roughness on top of the lines. The situation is sketched in figure 5.1. We will study vertical sidewalls for most of this chapter, but we will introduce a sidewall angle $\alpha < 90^\circ$ in section 5.6.1. A situation with roughness on top (but not on the sidewalls) will be studied in section 5.8. A more complete study should include top and bottom corner rounding, and have roughness everywhere on the structure.

We then use the simulator from chapters 2 and 3 to obtain the corresponding SEM image. Contours are extracted from this image using a thresholding method (section 5.2.2). The resulting measured power spectrum can then be compared to the known input PSD.

Our lines are 1024 nm long, with a grid spacing of $\Delta y = 1$ nm in the direction along the line. The correlation length in this direction is usually $\xi = 10$ nm unless mentioned otherwise. The choice of Δy coincides with the pixel size of the SEM images, which is 1 nm in both x and y directions. This was done to avoid any aliasing effects. In the vertical direction, where we sometimes choose a smaller correlation length, we used a grid spacing of $\Delta z = 0.5$ nm.

The resist lines are made of PMMA, the substrate is silicon. Assuming identical roughness parameters, we have performed the same study with silicon lines on a silicon substrate. The results were very similar. This gives confidence that the results are general: they are dominated by the geometry of the sample and not material contrast. In reality, the roughness of resist (after development) and silicon (after etch) features are very different, but it appears that the SEM response to roughness is not very material-dependent.

5.2.1 Generation of rough features

We assume that sidewall roughness can be characterised as stationary gaussian roughness. Such roughness can be quantified with the autocorrelation or autocovariance function. A common choice for the autocovariance function is a stretched exponential [134],

$$R(x) = \sigma^2 \exp\left(-\left|\frac{x}{\xi}\right|^{2\alpha}\right), \quad (5.1)$$

where ξ is the correlation length and α is known as the roughness (or Hurst) exponent. Instead of this stretched exponential, we will use the model of Palasantzas [135], which corresponds to the following autocovariance function:

$$R(x) = \frac{2^{1-H} \sigma^2}{\Gamma(H)} \left|\frac{x}{\xi}\right|^H K_H\left(\left|\frac{x}{\xi}\right|\right). \quad (5.2)$$

$K_H(x)$ is the modified Bessel function of the second kind. If $H = 1/2$, this is identical to the autocovariance of equation (5.1) with $\alpha = 1/2$. For $H \neq 1/2$, equations (5.1) and (5.2) remain similar, with H taking the role of α . The main advantage of Palasantzas's choice is that the PSD has a known analytical form for all roughness exponents.

The PSD in one dimension can be obtained by Fourier transformation of

equation (5.2):

$$PSD(f) = \int_{-\infty}^{\infty} dx e^{-2\pi ifx} \frac{2^{1-H} \sigma^2}{\Gamma(H)} \left| \frac{x}{\xi} \right|^H K_H \left(\left| \frac{x}{\xi} \right| \right) \quad (5.3)$$

$$= \frac{\Gamma(H + 1/2)}{\Gamma(H)} \frac{2\sqrt{\pi} \xi \sigma^2}{[1 + (2\pi f \xi)^2]^{H+1/2}}. \quad (5.4)$$

f is the spatial frequency, $1/\text{wavelength}$. $\Gamma(x)$ is the gamma function.

The two-dimensional PSD can also be found from the autocovariance function of equation (5.2). To allow for different correlation lengths in the x and y directions, we replace $|x/\xi|$ by $\sqrt{(x/\xi_x)^2 + (y/\xi_y)^2}$:

$$R(x, y) = \frac{2^{1-H} \sigma^2}{\Gamma(H)} \left[\left(\frac{x}{\xi_x} \right)^2 + \left(\frac{y}{\xi_y} \right)^2 \right]^{H/2} K_H \left(\sqrt{\left(\frac{x}{\xi_x} \right)^2 + \left(\frac{y}{\xi_y} \right)^2} \right). \quad (5.5)$$

The 2D PSD can be obtained by 2D Fourier transformation:

$$PSD(f_x, f_y) = \int_{-\infty}^{\infty} dx dy e^{-2\pi i(f_x x + f_y y)} R(x, y) \quad (5.6)$$

$$= \frac{4\pi H \sigma^2 \xi_x \xi_y}{[1 + (2\pi f_x \xi_x)^2 + (2\pi f_y \xi_y)^2]^{H+1}}. \quad (5.7)$$

For $\xi_x = \xi_y$, this equation reduces to the original isotropic form in the work of Palasantzas [135].

We use the method of Thorsos [136] to generate random rough sidewalls. This method generates the Fourier transform of the sidewall by taking the amplitude from the desired PSD and randomizing the phase. The real-space sidewall can then be found by an inverse Fourier transform. As explained by Mack [137], the variance of the data produced by this method is biased compared to the intended value. Additionally, because the roughness is periodic, the shape of the PSD is biased if the length of the feature is on the order of the correlation length.

The bias is small when $\Delta x \ll \xi \ll L$, where Δx is the grid spacing, ξ is the correlation length of the roughness spectrum, and L is the total size of the domain. In the direction along the line, this requirement is automatically satisfied by our choices of Δx , ξ and L . In the vertical direction, however, it is often the case that ξ is close to L or larger. Our solution is to use the Thorsos method to generate much higher sidewalls than required, and using only a section of

the desired size. After this step, the generated PSD closely follows the intended profile. Any remaining bias between the target and true 3σ SWR is small, but we subtract the mean and scale the amplitudes to guarantee that the 3σ SWR is exactly as intended.

5.2.2 Contouring method

A large variety of contouring methods is available in the literature. The reason why this large variety exists is that a simple thresholding method is plagued by robustness issues due to SEM noise. Filtering solves the robustness issue, but filters along the direction of the line must be avoided because they inevitably also affect the LER signal.

Alternative methods based on fitting reference linescan models [138, 139] do not rely on filtering or a threshold, and work well even for very noisy images. However, such methods assume that the linescan—and with it, the sidewall's vertical profile—is the same for every position along the line. The applicability of these methods has not yet been proven for features with SWR, and therefore, we do not use such a method.

We opted for a simple thresholding method, with minimal filtering to make it robust. We have set the threshold to 60% between the minimum and maximal signal of every linescan. This value does not significantly influence our results, unless the feature has a sidewall angle or the trenches are narrower than the height of the lines. These situations are further discussed in section 5.6. We use a $\sigma = 1$ nm Gaussian blur in the direction perpendicular to the line. Since this filter does not operate in the direction along the line, it does not affect the measured PSD as strongly as a parallel filter. Using simulated images with very low noise as a benchmark, we have empirically determined that a $\sigma = 1$ nm Gaussian blur does not noticeably affect the measured PSD.

Despite this, the SEM signal in a linescan may cross the threshold multiple times. Only when the slope is as expected (*i.e.* positive for a left edge, negative for a right edge) do we consider a threshold crossing as a candidate for edge detection. In the exceedingly rare case that there are still multiple candidates, we select the one closest to the mean edge position.

Because this is a simulation study, we can make the noise levels in the SEM images arbitrarily low. We have used 100 electrons per pixel, disabled Poisson shot noise from the electron source, and excluded detector noise. This gives images that are less noisy than those from a real CD-SEM. Due to higher noise in real SEM images, and due to effects such as shrinking and charging, a more

sophisticated contouring algorithm may be required for real SEM images.

We now have a set of measured edge positions, one for each edge in the image. The PSDs for each of these edges can be determined and averaged. Though the contouring method always finds a reasonable edge position, it remains sensitive to statistical noise of the SEM. Assuming that the noise is uncorrelated from pixel to pixel, this manifests itself as a white “noise floor”, which can be subtracted from the measured PSD [122, 140].

If the noise floor is very low, the measured PSD is not always completely flat at high frequencies. We measure the height of the floor by fitting the function

$$PSD(f) = a f^{-b} e^{-cf^2} + d \quad (5.8)$$

on a log-scale to the highest half of the frequencies. The f^{-b} term represents the typical power-law behaviour at high frequencies. The e^{-cf^2} term accounts for possible additional quenching of high frequencies due to the finite width of the electron beam. d is the actual noise floor.

Though equation (5.8) makes a lot of assumptions about the shape of the PSD at high frequencies, we only use this procedure to measure d . The value of d is not very sensitive to the exact form of the fit function, but we chose equation (5.8) because it can be physically justified.

5.3 No vertical roughness

We first study the case without vertical roughness. We use the PSD of equation (5.4) with $\sigma = 1$ nm, $\xi = 10$ nm and $H = 0.5$. We study dense lines and spaces with a 32 nm half-pitch.

The measured PSD, for an infinitely narrow electron beam with 300 eV landing energy, is shown in figure 5.2a. It also shows the measurement of the noise floor using the fit of equation (5.8). Figure 5.2b shows the same PSD, with the noise floor subtracted. The measured PSD differs from the input PSD only in the very high frequencies. This divergence is most likely due to uncertainty in the noise floor measurement due to our choice of fitting function. We have run the same simulation with different electron beam landing energies (300 eV, 800 eV and 3000 eV), and with a smaller x-pixel size (0.5 nm). The results were practically indistinguishable from figure 5.2b. This implies that the high-frequency divergence in figure 5.2b is not due to the electron scattering cascade or material parameters.

In figures 5.2c and 5.2d, we use a more realistic Gaussian profile for the electron beam. The full width containing 50% of the current (FW50) is 3 nm, similar

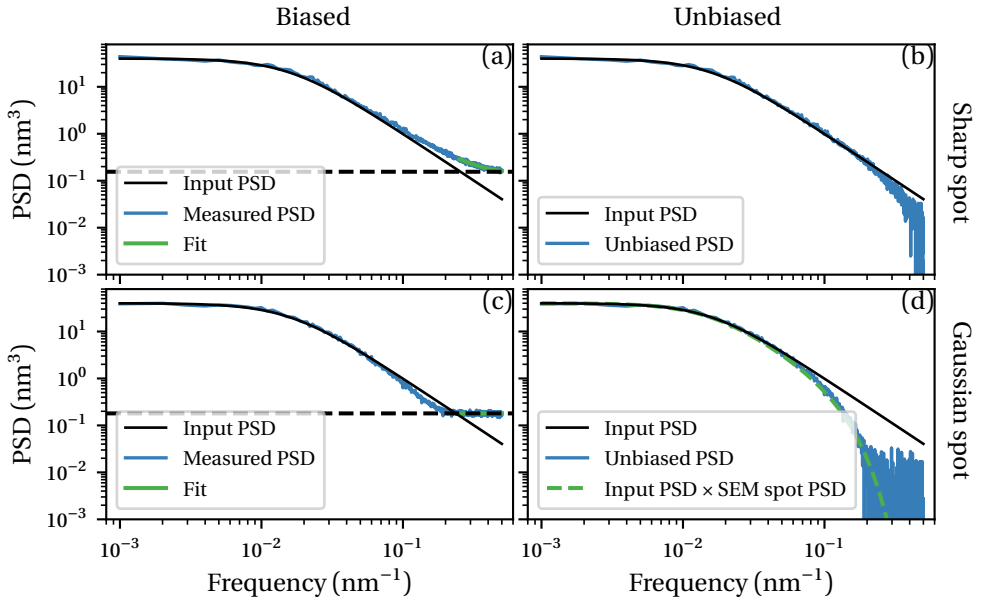


Figure 5.2: The input PSD (black curve) is compared to the PSD measured from simulated SEM images (blue curves). There is no roughness in the vertical direction. The figures on the left show the measurement of the noise floor, the figures on the right show the result after noise floor subtraction. In (a) and (b) the electron beam is infinitely sharp; in (c) and (d) the electron beam has a Gaussian profile with 3 nm FW50. The dashed curve in (d) indicates the input PSD convolved with the SEM spot.




Figure 5.3: Illustration of the projection model. An example 3D rough line is shown. The contour of every slice of the geometry is projected onto the plane above the feature (grey lines). The red lines follow the outermost extrusions of the grey lines. The projection model assumes that the SEM observes this red line.

to current CD-SEMs. The effect of this finite spot size is a suppression of the PSD at higher frequencies. This can be modelled by convolving the PSD of the line by the PSD of the spot. The blurring effect due to the finite spot size is a very dominant effect.

Because the effect of a finite electron beam spot size can be modelled well with a convolution, all future simulations in this work will be performed with an infinitely sharp beam.

5.4 Isolated lines

We now also add roughness in the vertical direction according to equation (5.7). For simplicity, we consider isolated lines first, before going to dense lines & spaces in the next section. The energy of the electron beam remains 300 eV. The effect of larger beam energies is investigated in section 5.7.

We will use the following simplistic model to understand the results. A typical SEM signal is bright when the beam lands near a sharp edge, because electrons may escape from the side. We may expect that, when the beam lands on an extrusion on the rough sidewall, the SEM signal is also very bright. This means the SEM contour will go around all extrusions as seen from the top. We call this the “projection model”, which is illustrated in figure 5.3.

Numerically, the projection model can be quantified by generating a random 3D rough sidewall. Then, at each position along the line, the maximum excursion in the vertical direction is taken. In this paper, we repeat this procedure for a very large number of sidewalls and find the average PSD. We have not

been able to find an analytical expression.

The projection model gives a vertical averaging effect. If the correlation length in the vertical direction, ξ_z , is very large, the structure looks like there is no roughness in the vertical direction. Hence, the measured PSD will be similar to the PSD of a single slice of the structure. If ξ_z is very small, the projection model averages over many correlation lengths in the vertical direction. Hence, the observed mean edge position will shift outwards, and the roughness will be less than that of a single slice. The amount of averaging depends on the number of correlation lengths in the vertical direction: scaling both the height of the feature and ξ_z by the same factor gives the same final result.

A typical SEM image of an isolated 3D rough line is shown in figure 5.4a. This figure also shows the contours detected by the contouring algorithm, the true contour of the top slice of the line, and the projection model's contour. By eye, it is already clear that the projection model predicts the SEM contour quite well. Figure 5.4b shows a selected scanline from figure 5.4a. This confirms the effects that lead us to hypothesise the projection model. The SEM signal already brightens for $x < 16$ nm (where the top of the line is), which is a well-known effect. The signal remains bright in the range $16 \text{ nm} < x < 19$ nm, where the electron beam lands on the various extrusions of the sidewall. The edge detected by the contouring algorithm is therefore close to the outermost extrusion.

The accuracy of the projection model can be verified by comparing its predicted PSD to the measurement by the SEM. Such a comparison is shown in figure 5.5. The agreement between the projection model and the SEM measurement is excellent. The SEM measurement slightly underestimates the signal at very high frequencies, similarly to the situation of figure 5.2.

It is clear from figure 5.5 that the measured 3σ LER (given by the area below the green curve), is biased significantly compared to the true 3σ SWR (given by the area below the black curve). This is consistent with the simulation and experimental studies mentioned in the introduction [91, 124, 129–133].

Figure 5.6 shows the bias of measured 3σ values for a range of ξ_z/h , where h is the height of the feature. We have included the effect of the CD-SEM's spot size in this figure. The full SEM simulations were done with a Gaussian electron beam spot. The projection model was convolved with the appropriate spot PSD (as in figure 5.2). The agreement between the projection model and the full simulations is quite satisfactory.

Interpreted differently, figure 5.6 predicts that measured LER scales with film thickness if the on-wafer SWR is constant. This is a pure metrology effect. It is consistent with previous experimental observations [141, 142], where LER




Figure 5.4: (a) Typical SEM image of an isolated 3D rough line. The contour detected by the contouring algorithm (red curve), as well as the contour of the top slice (purple) and the projection model (blue) are overlaid. The yellow horizontal line marks a selected scanline, which is reproduced in (b). The black curve represents the grayscale value of the SEM image. The positions of the top, projection model and SEM contour are shown as vertical lines. This is an 80 nm high line with $\sigma = 1$ nm, $\xi_y = \xi_z = 10$ nm, and $H = 0.5$.




Figure 5.5: The input PSD of a single slice in the sidewall (black curve) is compared to the PSD predicted by the projection model (blue) and the PSD measured by SEM (green). This figure was made for the same parameters as figure 5.4.




Figure 5.6: Comparing the bias in 3σ LER predicted by the projection model (solid curve) to observation from SEM images for a range of heights (points). The dashed line indicates the true 3σ SWR. This figure was made for $\xi_y = 10$ nm and $H = 0.5$. A 3 nm FW50 Gaussian spot for the CD-SEM is included in these simulations.

is seen to increase as resist thicknesses are reduced. It is possible that a similar metrology component plays a role in those experimental studies, in addition to any change in on-wafer roughness.

As shown by Verduin *et al.* [91], the correlation length measured by CD-SEM is biased if there is roughness in the vertical direction. We quantify this effect in figure 5.7a. According to the projection model, only the ratio $\xi_{\text{measured}}/\xi_y$ is biased: scaling ξ_y results in a scaling of the PSD to lower or higher frequencies, while the shape is preserved. The bias in measured correlation length shows an interesting trend. If there are many correlation lengths in the vertical direction (small ξ_z/h), the measured correlation length is significantly smaller than the true correlation length in the horizontal direction. When the structure is approximately two correlation lengths high, the measured correlation length peaks at a larger value than the true correlation length in the horizontal direction. For large ξ_z/h , the situation becomes similar to the case when there is no vertical roughness, and the bias disappears.

It can be seen in figure 5.5 that the projection model gives a lower Hurst exponent than the input PSD. We quantify this bias in figure 5.7b. The bias is almost linear, but it depends on ξ_z/h . As ξ_z/h approaches infinity, the absence of vertical roughness should make the measured H equal to the input value. The curves in figure 5.7b seem to agree with that trend.

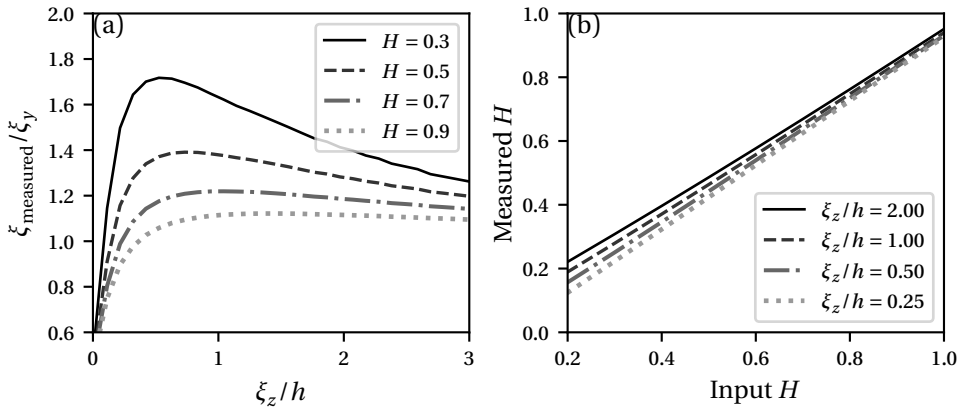


Figure 5.7: The bias in measured correlation length and Hurst exponent, predicted by the projection model. The horizontal axis in figure (a) indicates the number of correlation lengths in the vertical direction; the vertical axis shows the measured correlation length relative to the true correlation length along the line. The horizontal axis in figure (b) indicates the true Hurst exponent; the vertical axis shows the measured Hurst exponent. The lines are predictions from the projection model, full SEM simulations are not shown.

We conclude this section by restating that the projection model predicts the SEM contour and its PSD very well for isolated lines. The projection model follows from purely geometrical arguments. This makes it very simple to understand and quantify. We did not attempt to find an analytical expression for the projection model PSD, but it is computationally very cheap to find numerical approximations by brute-force Monte Carlo simulation. This enables building a large library to translate “projected PSDs” back to the true sidewall PSD.

5.5 Dense lines and spaces

We are now interested in the more common situation of a dense pattern with 50% lines and 50% spaces. An example SEM image is shown in figure 5.8.

Comparing figure 5.8 to figure 5.4, we observe that some parts of the sidewall are darker in the dense line-space pattern. It is likely that electrons escaping from lower layers of the structure are blocked by the neighbouring walls. As a result, those deeper layers show up darker in the SEM image. Parts of the image that are darker than the SEM contouring threshold are not considered part of the line. Therefore, the projection model does not predict the SEM contour as

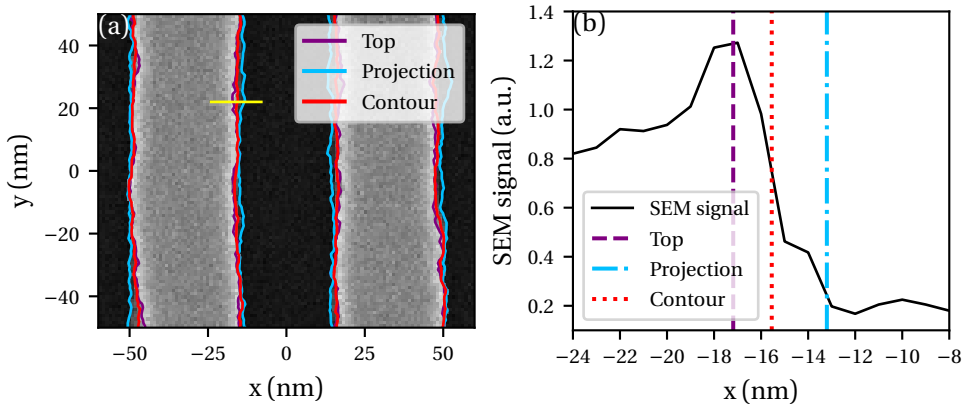


Figure 5.8: Same as figure 5.4, with trenches between the lines. (a) shows a SEM image with the true positions of the top slice and projection model overlaid, as well as the contour detected using the image. (b) shows a selected line of the image, marked by the horizontal yellow line in (a). The resist lines are 80 nm high, trenches are 32 nm wide.

well as it did for the case of purely isolated lines.

We will now attempt to model this effect. Because the SEM image gets progressively darker for deeper layers, we may hypothesise that the contouring threshold for the SEM image translates to a threshold depth in the actual resist pattern. We may expect that all sidewall features above this depth become part of the SEM contour. All sidewall extrusions below this depth are too dark, and the contouring algorithm considers them part of the trench. Therefore, we will investigate a “cutoff projection model”, in which we assume that the SEM contour follows the projection model above a certain cutoff depth, which we need to calibrate. A typical result is shown in figure 5.9. This result was obtained by tuning the projection model’s cutoff depth such that its low-frequency $PSD(0)$ matches the measured $PSD(0)$.

The match between model and measured PSDs in figure 5.9b is much worse than in figures 5.2 and 5.5. High frequencies are much more suppressed in the SEM image than the simple cutoff projection model predicts. The same is also clear from the real-space figure 5.9a, where the cutoff projection model does not follow the SEM contour very closely despite our efforts to match the PSDs. This indicates that the cutoff projection model is too simplistic. This is unfortunate, but not very surprising. There are other effects at play in the image formation of a SEM. Sharp geometric contrast generally makes the image bright,




Figure 5.9: (a) Same as figure 5.8a, but the light blue curve now shows the projection model with a cutoff at 30 nm depth. (b) Corresponding PSDs. The 30 nm cutoff depth was chosen to match the measured $PSD(0)$.

which is why the projection model works well. But the extrusions on the sidewall, which are of the order $\sigma = 1$ nm large in one direction, are smaller than the interaction volume of the incoming electrons. Therefore, it is likely that the exact brightness in the image depends on the local geometry of the sidewall and not just the depth. Some features above our chosen cutoff depth may therefore be darker than the contouring threshold, while some features below the cutoff depth may be brighter. The relation between the measured PSD and the contouring threshold is discussed in section 5.6.

It is worth investigating under what conditions lines can be considered “isolated enough” for the projection model (without cutoff) to be applicable. Figure 5.10 shows how the measured 3σ values evolve as trenches become wider. When the trench is wider than the feature height, the projection model predicts the measured 3σ value quite well. The projection model becomes worse as the trenches get narrower than the line height.

5.6 Contouring threshold

Throughout this paper, we have used a contouring threshold of 60% between the minimum and maximum of the SEM signal. This is a common choice in industry. Most of our results are not sensitive to the choice of contouring threshold, except two situations.




Figure 5.10: The bias in measured 3σ LER as a function of trench width. These simulations were done for 40 nm high lines, with $\xi_y = \xi_z = 10$ nm and $H = 0.5$.

5.6.1 Sidewall angle

When there is a sidewall angle, the projection model works well for isolated lines. One thing to note is that the contouring threshold should not be set too high.

Figure 5.11a shows a SEM linescan over a trapezoidal line edge without any roughness. Figure 5.11b shows a SEM image of a line with a rough sidewall (as illustrated in figure 5.1). Setting the contouring threshold too high puts the contour on top of the sidewall rather than near the outermost edge. This also makes it very sensitive to noise. High contouring thresholds should therefore be avoided: not because the projection model fails, but also because the contouring algorithm becomes too sensitive to noise.

After doing many simulations, we conclude that the projection model works well when the contouring threshold is 60% or less. Higher thresholds than 60% are very uncommon in literature. For practical purposes, we conclude that the projection model works not just for vertical sidewalls, but also when the line has a sidewall angle.

5.6.2 Dense lines and spaces

As discussed, lower layers of dense line-space patterns are darker in SEM images due to a shadowing effect from neighbouring lines. This effect can be seen in figure 5.8. By eye, it is clear that a lower contouring threshold will put the SEM contour closer to the projection model. An obvious question now is whether




Figure 5.11: SEM images of isolated trapezoidal structures, 40 nm high and with an 85° sidewall angle. (a) SEM linescan over a structure without any roughness. The shaded area represents the shape of the structure. (b) SEM image of an isolated 3D rough line. The red and green curves show the measured contour lines at different threshold settings. The roughness parameters are the same as in figure 5.4.

the projection model becomes more applicable as the contouring threshold is lowered, and, if so, whether it is possible to find a threshold setting for which the projection model can be universally trusted.

Typical PSDs are shown in figure 5.12 for various contouring thresholds. As the contouring threshold is lowered, the measured PSD becomes more similar to the projection model: the PSD at low frequencies decreases while the PSD at high frequencies increases. However, even at a very low 20% contouring threshold, the difference remains obvious. It is conceivable that an even lower contouring threshold will eventually match the projection model, but judging from the trend in figure 5.12, this looks unlikely. Even if there is a very low threshold for which the projection model is matched, this will not be a practically useful threshold because such extreme contouring thresholds are very sensitive to SEM noise. Therefore, we conclude that there is no practical threshold setting at which the projection model becomes applicable for these very dense lines and spaces.




Figure 5.12: Measured PSDs for various contouring thresholds in a dense lines-space pattern. Physical parameters are the same as in figure 5.9.

5.7 Energy dependence

The simulations in the previous sections were all done for 300 eV electron beams. The projection model, which follows from purely geometric arguments, works very well for this low landing energy. Higher landing energies lead to a larger “interaction volume” in which the electrons scatter. The result is an effective blur of the SEM image.

This effect is demonstrated in figure 5.13. It shows that increasing the landing energy leads to a suppression of the measured power at high frequencies. This effect is already significant at 1 keV beam energy.

5.8 Top roughness

Resist features have roughness everywhere, not only along the sidewall. The projection model predicts that roughness on the top of the resist does not influence the measured LER. It is good to verify this.

We study resist lines with roughness only on the top, as illustrated in figure 5.14. The sidewalls are now flat, the sidewall angle is 90° . If the projection model is correct, the SEM contour is pure white noise. The roughness on top of the feature is described by the same PSD as before, equation (5.7), with $\sigma = 1$ nm, $\xi_x = \xi_y = 10$ nm and $H = 0.5$.

Simulation results are shown in figure 5.15. The figures on the left show con-




Figure 5.13: The measured PSD for different landing energies of the electron beam. All other parameters are the same as in figures 5.4 and 5.5.




Figure 5.14: Sketch of the line profile studied in section 5.8. We study lines with perfectly straight and vertical sidewalls. The top of the resist is modulated by a roughness model.

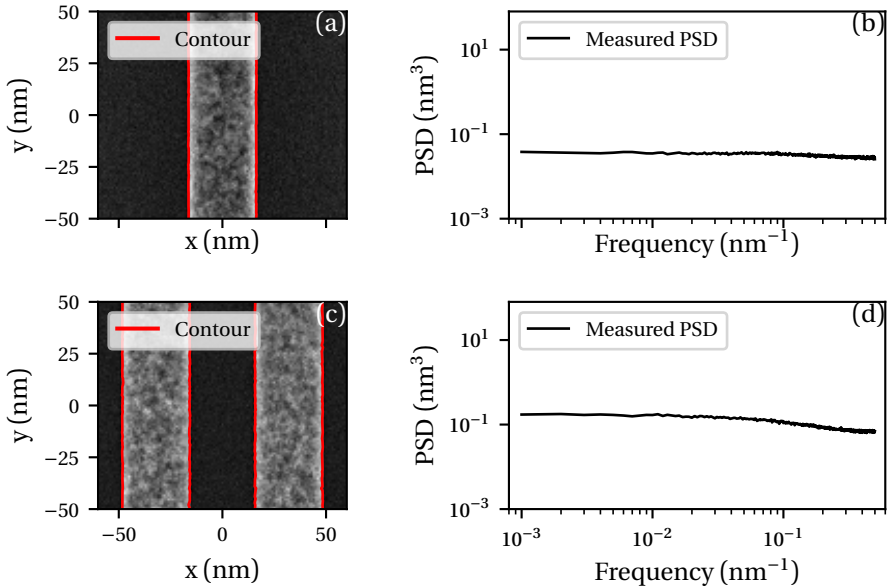


Figure 5.15: Simulated SEM images and measured PSDs for features with roughness on top of the lines only, as illustrated in figure 5.14. (a) and (b) are for isolated lines; (c) and (d) for dense lines & spaces. The noise floor was not subtracted from the measured PSDs in figures (b) and (d).

toured SEM images, the figures on the right show the measured PSDs. The noise floor was not subtracted from the measured PSDs. The measured PSDs are almost flat, but not perfectly.

We may propose the following explanation. The top roughness is clearly visible in the SEM images. Because of this, the brightness of the line edges is also modulated. The brightness of the substrate is (on average) constant. The bright parts of the edge push the measured contour slightly outwards, hence the top roughness “leaks” into the measured LER. However, this effect is negligible compared to the usual power contained in LER signals (*cf.* figures 5.2, 5.5, 5.9). We estimate that the total variance due to non-white noise in figure 5.15b is 0.0014 nm^2 . In figure 5.15d, this is 0.011 nm^2 .

5.9 Conclusions

We have studied the influence of SWR on LER measured by top-down CD-SEM. This was done by generating features with known SWR, simulating correspond-

ing CD-SEM images and measuring LER from those images.

We have verified that the popular method of PSD analysis works very well if the structures have no roughness in the vertical direction. The dominant factor limiting the measurement of the true on-wafer PSD is the electron beam's spot size. This causes a blurring effect, suppressing the PSD at high spatial frequencies.

If the structures are isolated and rough in the vertical direction, the CD-SEM observes the outermost extrusion over the full height of the structure. This is the case for PMMA lines on a silicon substrate as shown here, but we have also verified this for silicon lines on a silicon substrate. This simple geometrical interpretation is therefore not due to material contrast, and it is not sensitive to details of the electron scattering process. It is therefore likely that this effect exists for arbitrary feature shapes. We have verified this explicitly for trapezoids with a sidewall angle and lines with roughness on the top.

As a result of this projection effect, the PSD measured by the SEM is biased significantly. Especially the 3σ LER, which is often considered the most important quantity, suffers heavily from this bias.

For dense lines and spaces, the situation is complicated by the fact that the SEM is less sensitive to the lower layers. As a result, the measured PSD lies between the PSD of the top slice and that of the projection model. Empirically, we have found that the projection model works well as long as the trenches between the lines are wider than the line height.

The ultimate goal of roughness metrology is usually to predict device performance or statistical defects, not to measure the true SWR of resist patterns. An interesting question is whether subsequent process steps (such as etch) undergo a similar projection effect as top-down CD-SEM measurements. If that turns out to be the case, a LER measurement by CD-SEM may be a good predictor of device performance. However, this work shows that care must be taken for dense line-space patterns (and most likely contact holes), because the SEM is most sensitive to upper layers, which is unlikely to be the case for etch. Also, if CD-SEM is used for benchmarking and improving resists, the full sidewall roughness is an interesting quantity. We have shown that measured LER depends on film thickness, highlighting the importance of benchmarking resists under identical circumstances.

Chapter 6

What is missing in our simulations

Chapter 2 summarizes the current state of the art of electron-matter interaction. In this chapter, we will revisit some of the major assumptions and approximations which are present in the theory. In sections 6.1–6.7, we will discuss these theoretical aspects and assess their importance. We will then discuss why practical experiments and theory will never be in perfect agreement, even if the theory is completely correct, in section 6.8.

It is not our goal to *solve* any of the issues presented in this chapter. Any important effects that can be easily addressed are already part of our modelling. This chapter lists inherent model uncertainties, effects that we do not know how to model, and effects that have been modelled in literature but represent only small corrections. While an effort has been made to present a comprehensive list in this chapter, it is possible that some effects have been forgotten.

Before getting into the model details, it is worth emphasising one fundamental issue with this type of simulator. It addresses the problem in the wrong direction, as we are fundamentally more interested in the inverse problem. In reality, one has a sample of interest but doesn't know what it looks like, and uses an electron microscope to find out. The simulator works in a similar way: even with perfect physical models, one must assume a full three-dimensional geometry before the corresponding electron microscope image can be simulated. It is unlikely that the sample's full three-dimensional structural information can be uniquely captured in a few two-dimensional images¹, hence we should not expect to be able to make a model that can recover the geometry from an experimental image.

If the sample is almost known, one may vary a limited number of unknown

¹One can vary the electron beam energy, and obtain two images (SE and BSE) per exposure.

parameters. This has been done before [17, 143], but such studies must deal with uncertainties from three categories: imperfect modelling, imperfect assumptions about the sample, and imperfect experimental conditions.

6.1 Optical data model

We will start with the optical-data model, which we use to describe inelastic scattering. The mean free path is given by equation (2.34), which reads (in non-relativistic form)

$$\lambda^{-1} = \frac{1}{\pi a_0 E} \int d\omega \int \frac{dq}{q} \operatorname{Im} \left[-\frac{1}{\epsilon(q, \omega)} \right]. \quad (6.1)$$

Here, a_0 is the Bohr radius, E is the electron's kinetic energy, and q and ω represent the transferred momentum and energy, respectively.

6.1.1 First Born approximation

In deriving equation (6.1), it was assumed that the material responds linearly to the electron's electric field:

$$\mathbf{D}(\mathbf{q}, \omega) = \epsilon_0 \epsilon(\mathbf{q}, \omega) \mathbf{E}(\mathbf{q}, \omega). \quad (6.2)$$

Here, $\epsilon(\mathbf{q}, \omega)$ is the material's dielectric function and $\mathbf{E}(\mathbf{q}, \omega)$ is the “bare” electric field from the electron.

Formally, this linearity assumption is equivalent to a first-order Born approximation [21]. Physically, the first Born approximation assumes that the incoming electron's wave function is not substantially altered by the scattering potential. This means that multiple scattering events in the same potential are ignored, as well as the constant acceleration and deceleration the electron experiences as it moves between atoms in the solid. As a rough guide, the first Born approximation is valid when the electron's energy is much greater than the potential in which it is scattered, and it may break down at lower energies. There is some debate in literature about the exact point where it starts to break down. According to Lindhard [36, page 42], the Born approximation should be valid down to the Rydberg energy (13.6 eV). He also gives two reasons why the Born approximation may still work at even lower energies. First, the valence electrons in the material screen the bare charge of the electron, so the effective interaction is less strong than the bare Coulomb interaction between electrons.

Second, the relevant quantity for interactions is the relative velocity between the incident and valence electrons, which is at least on the order of the Fermi velocity. Ritchie [21] further comments that the Born approximation is exact in the case of Coulomb scattering, so we may expect it to be reasonable here as well. A different lower bound is given by Tung *et al.* [23] and Mott and Massey [144]. They say that, in practice, the Born approximation gives acceptably accurate results when the energy of the electron exceeds seven times the Fermi energy.

Given the discrepancy in literature, it is now time to see what the effect of a second-order Born approximation would be. Tung *et al.* [23] give a mathematical expression for the second order Born approximation. Figure 6.1 shows the differential (with respect to energy loss) inverse mean free path, $d\lambda^{-1}/d\omega$, for two energies in silicon. There are two curves: with and without the second-order Born correction taken into account. The second-order Born correction is negligible for a 500 eV electron, but has an appreciable impact at 100 eV. The correction from the paper of Tung *et al.* [23], allows the final mean free path to be negative. No mention is made of this in reference [23], but with the Penn optical-data model and silicon as material, this turns out to happen at $\lesssim 60$ eV. Such behaviour is clearly unphysical, but it should be possible to correct for this by taking higher orders of the Born series into account. Figure 6.2 shows the impact on the mean free path.

6.1.2 Exchange effects

The dielectric function model, equation (6.1), does not account for exchange effects between the incident and target electrons. At least two ways of taking the exchange effect into account have been proposed.

One was proposed by Tung *et al.* [145]. They recognise that, in the limit of high energy and high energy loss, equation (6.1) reduces to the non-exchange part of the Møller electron–electron scattering cross section. They add two additional terms, similar to the right-hand side of equation (6.1), such that the resulting expression reduces to the exchange-corrected Møller cross section in the high-energy limit. This approach is also used in the optical-data model of Ashley [43], and by Tung *et al.* [23].

Ochkur [146, 147] performed a real quantum-mechanical calculation. He uses the Born-Oppenheimer approximation, expanding it in terms of $1/E$, and using the leading term. This approximate result is different from that of Tung *et al.* [145]. Both approximations have a qualitatively similar effect: the scattering




Figure 6.1: The differential inverse mean free path (DIMFP), $d\lambda^{-1} / d\omega$, for inelastic scattering in silicon. (a) is for a 100 eV electron, (b) is for 500 eV. The horizontal axis indicates the energy loss, the vertical axis is proportional to the probability. Two curves are shown, for the the first (solid) and second (dashed) order Born correction. In (b), the curves are on top of each other.




Figure 6.2: The inverse mean free path of an electron in silicon, for the first (solid curve) and second-order (dashed) Born approximation. The mean free path for the second-order Born approximation is cut off at 60 eV because the differential inverse mean free path becomes unphysical below this energy.

rate is reduced, with the effect being strongest at low electron energies. Ochkur's approximation is used by Fernández-Varea *et al.* [148].

The effect of both exchange corrections on the differential inverse mean free path can be seen in figure 6.3. Similarly to the second-order Born correction, the exchange effect is negligible for 500 eV electrons but is more significant at lower energies. The corresponding mean free paths are shown in figure 6.4. The exchange corrections of Tung and Ochkur are qualitatively quite similar over the entire energy range. The exchange-corrected mean free path is about a factor 2 higher around 35 eV, but the difference is usually less. As we saw in section 3.6, a factor 2 can have a considerable effect on electron yields.

6.1.3 Transverse component of electromagnetic field

The derivation of equation (6.1) assumes that the electron is non-relativistic. This allowed us to approximate the problem as an electrostatic one: we worked in the Coulomb gauge with the vector potential set to zero. This approximation becomes invalid for relativistic electrons.

An expression including the transverse component is given by Fernández-Varea *et al.* [149]. In a different article [150], authors from the same group state that this can be safely neglected for particles with kinetic energies less than 0.5 MeV: much more than the energies we consider in this thesis.

6.1.4 Dielectric function is a tensor

We have not only assumed a linear relationship between the electric field and the displacement field, equation (6.2), but we have also assumed that the dielectric function is a scalar. This corresponds to the assumption of an isotropic medium. In general, the dielectric function $\epsilon(\mathbf{q}, \omega)$ can be a tensor. This is most likely acceptable when the material is polycrystalline or amorphous, but it may be problematic in crystalline solids.

The generalised form of equation (6.1) that allows for a tensorial dielectric function is given by Saslow and Reiter [151]. In practice, the full dielectric tensor over a large energy range is often difficult to find. For the present discussion, we shall limit ourselves to some typical values for interesting materials. These are given in table 6.1. It seems that the anisotropy is often on the order of 10% or less.




Figure 6.3: The differential inverse mean free path (DIMFP), $d\lambda^{-1} / d\omega$, for inelastic scattering in silicon. (a) is for a 100 eV electron, (b) is for 500 eV. The horizontal axis indicates the energy loss, the vertical axis is proportional to the probability. Three curves are shown: a model without exchange correction (solid curve), with the exchange correction of Tung *et al.* [145] (dashed) and with that of Ochkur [146, 147] (dash-dotted). In (b), the curves are on top of each other




Figure 6.4: The inelastic mean free path of an electron in silicon, without exchange correction (solid curve), with that of Tung *et al.* [145] (dashed) and with that of Ochkur [146, 147] (dash-dotted).

Table 6.1: Static (ϵ^0) and high-frequency (ϵ^∞) values for the dielectric tensor. These values were calculated using density functional theory by Sevik and Bulutay [152].

	$\epsilon_{xx}^0 = \epsilon_{yy}^0$	ϵ_{zz}^0	$\epsilon_{xx}^\infty = \epsilon_{yy}^\infty$	ϵ_{zz}^∞
SiO ₂ (α -quartz)	4.643	4.847	2.514	2.545
SiO ₂ (α -cristobalite)	4.140	3.938	2.274	2.264
SiO ₂ (β -cristobalite)	3.770	3.770	2.078	2.078
SiO ₂ (stishovite)	10.877	8.645	3.341	3.510
GeO ₂ (α -quartz)	5.424	5.608	2.864	2.947
GeO ₂ (rutile)	10.876	8.747	3.679	3.945
α -Al ₂ O ₃	10.372	10.372	3.188	3.188
β -Si ₃ N ₄	8.053	8.053	4.211	4.294
β -Ge ₃ N ₄	8.702	8.643	4.558	4.667

6.1.5 The choice of model dielectric function

As discussed in section 2.4.4, we use a measured dielectric function in the optical ($q \approx 0$) regime, which is extended to $q > 0$ by fitting a model dielectric function. Each of these model dielectric functions comes with its own assumptions and approximations. Most of these models are in good agreement when the electron energy exceeds 100 eV, but their predictions diverge at lower energies. The details are discussed in section 2.4.4 and summarised in figure 2.8.

6.2 Inner-shell excitation

We use a model for cold, isolated, neutral atoms in the gas phase to describe inner-shell excitation. The calculations are from the LLNL Evaluated Electron Data Library (EEDL) [59], released as part of the ENDF/B library [60] in 2018.

6.2.1 Intrinsic model uncertainties

The calculations underpinning EEDL have intrinsic uncertainties. The documentation states that the uncertainty is largest for low electron energies (near the ionization threshold) and large atomic numbers.

Unfortunately, no explicit uncertainties are given. We can get a feel for the uncertainty by comparing the calculated ionization cross sections in EEDL to experimental measurement data. Figure 6.5 shows such a comparison for the K shell of silicon ($Z = 14$). This shows good agreement between computed and

experimental values, though the experimental values systematically appear to be slightly higher. Figure 6.6 shows the same for the K and L shells of gold. Here, the difference is significant: it can be a factor 2–3 in the energy range considered here. EEDL comes closer to experimental data for electron energies higher than those in figure 6.6, but those energies are not very interesting for us.

6.2.2 Solid-state effects

We use data for isolated neutral atoms to describe the inner shells of atoms in a solid. For light elements, which have a clear separation between deep electrons and valence electrons, this is unlikely to be a big problem. Valence electrons are perturbed in the solid state, forming bonds between the atoms, but the inner electrons are barely affected.

Let us take gold as an example: in its ground state, it has one electron in its only unfilled (6s) shell. It seems that gold has one valence electron, and indeed, compounds such as AuF and AuCl exist. However, because the energy differences between the outer shells are so low, it is quite easy for electrons to move to different shells, and form compounds such as AuF₃ and AuCl₃. A second point is that these heavier elements also have “intermediate shells,” which are too tightly bound to the core to form bonds, but still interact with other shells. To illustrate this point, figure 6.7 shows the energy-loss function of gold with the binding energies of all shells. At > 2 keV binding energies, the inner shells show up as sharp ridges in the energy-loss function, as expected for tightly-bound electrons. There are also several “medium-bound” shells, whose contributions to the energy-loss function appear as broad peaks. This indicates that these shells “mix” with others to a large extent. It is possible that these shells are also perturbed if the atom is part of a solid.

6.2.3 Free-electron gas model in the dielectric function

Inner shells are included in the dielectric function. This means that inner shells are treated using optical data (photoionization) which is extended to nonzero momentum transfers by a free-electron gas model (specifically, the Penn algorithm [42]). Fortunately, small momentum transfers are the most important thanks to (a) the factor $1/q$ in equation (6.1) and (b) the fact that $\text{Im}[-1/\epsilon]$ generally decreases for increasing q . It would, however, be more appropriate to consider the valence electrons and the inner-shell electrons separately.

We now introduce the following model, which allows us to consider valence and inner-shell electrons separately. We start with the energy-loss function of




Figure 6.5: Measured (points) and calculated (solid line) electron ionization cross sections of the K shell in silicon. Measurement uncertainties are as reported in the publications.




Figure 6.6: Measured (points) and calculated (solid lines) electron ionization cross sections of several subshells in gold, specifically the (a) K, (b) L1, (c) L2 and (d) L3 shells. Measurement uncertainties are as reported in the publications, the error bars are often smaller than the markers.




Figure 6.7: The energy-loss function of gold (solid curve), together with binding energies of all electron shells tabulated in EEDL (dashed vertical lines). The energy-loss function was taken from [30, 31].

silicon, and “subtract” the inner shells by fitting a ω^{-3} power law to the optical energy-loss function. The inverse third power is chosen because this is the large- ω limit of the free-electron gas model in equation (2.62). The result can be seen in figure 6.8. Then we add the inner shells separately, using the cross sections tabulated in EEDL.

The result, for silicon, is summarised in figure 6.9. Figure 6.9a shows three mean free paths: that in our current model (dash-dotted curve), that when we consider only the valence electrons with a free-electron gas model (solid curve), and that of inner-shell ionization separately. The mean free path of inner-shell ionization is much longer than that of valence electron excitation, meaning that it is a much less likely type of event. Because inner-shell excitation typically involves much larger energy transfers than valence excitation, it does have an important contribution to the overall stopping power. This can be seen in figure 6.9b, which shows the stopping power of the same three models.

If the Penn model described inner-shell excitation perfectly, we would have

$$\lambda_{\text{simulator}}^{-1} = \lambda_{\text{core}}^{-1} + \lambda_{\text{valence}}^{-1} \quad (6.3)$$

$$S_{\text{simulator}} = S_{\text{core}} + S_{\text{valence}}, \quad (6.4)$$

where λ^{-1} is the inverse mean free path and S represents stopping power. Figures 6.9c and 6.9d indicate how much deviation there is by plotting the ratios $\lambda_{\text{simulator}}^{-1}/(\lambda_{\text{core}}^{-1} + \lambda_{\text{valence}}^{-1}) - 1$ and $S_{\text{simulator}}/(S_{\text{core}} + S_{\text{valence}}) - 1$, respectively. We can see that the Penn model gives the correct inner-shell mean free path to




Figure 6.8: The energy-loss function of silicon, with (solid curve) and without (dashed) inner shells included.




Figure 6.9: (a) mean free path (MFP) and (b) stopping power (SP) in silicon. Solid curves: Penn model with only valence electrons (dashed curve in figure 6.8). Dashed curves: inner-shell ionization derived from EEDL. Dash-dotted curves: current simulator, using the Penn model for all electrons (solid curve in figure 6.8). (c) shows the discrepancy between the inverse MFPs $\lambda_{\text{simulator}}^{-1}$ and $\lambda_{\text{core}}^{-1} + \lambda_{\text{valence}}^{-1}$. (d) shows the same for stopping power.

within a few percent. However, the stopping power is significantly overestimated, by up to approximately 20%. This implies that the average energy loss per event is not described very well by the Penn algorithm.

We remind the reader of the current method we use for inner-shell excitation, as described in section 2.5.3. The mean free path is determined from the Penn model with inner shells included. Then, we choose whether or not ionization occurs by considering the relative ionization cross sections of valence electrons and inner-shell electrons as tabulated in EEDL. According to figure 6.9c, it may also be appropriate to use the absolute cross sections for inner shells from EEDL here, ignoring the valence electrons in EEDL. Then, we use the energy loss from the Penn model. According to figure 6.9d, it may be better to use the energy loss distribution from EEDL instead.

6.3 Low-energy elastic scattering

We use a combination of Mott scattering (at > 200 eV) and electron–acoustic phonon scattering (< 200 eV). These are described in detail in sections 2.6 and 2.7.

It is well known that electron–acoustic phonon scattering describes conduction phenomena very well. At the Fermi level, the contribution of Mott scattering must therefore be negligible. In particular, it must be less than what ELSEPA predicts. Conversely, it is also well established that Mott scattering provides an accurate description in the high-energy case. The phonon mean free path must therefore become negligible at these energies, but in practice that does not happen.

For this reason, it makes sense to use a phonon model at low energies, then smoothly “turn it off” at intermediate energies and transition to Mott scattering at high energies. Unfortunately, it is unclear where this transition region should lie. There is also no clear physical reason why electron–phonon scattering should “turn off” at the same rate as Mott scattering should “turn on”, but to avoid the extra complication, we are now going to assume that this is the case anyway. The inherent model uncertainty in Mott scattering becomes large around 100–200 eV, see figure 2.10. This is one of the reasons why Kieft and Bosch [10] chose this energy region to transition between the two models. The other argument was that at 100 eV, the electron’s De Broglie wavelength, $\lambda = h/\sqrt{2mE}$, is on the order of the interatomic distance. Typical lattice constants are on the order of 5 Å, and the shortest interatomic distance can be a bit less than that. Using the “vacuum” electron mass (not the effective mass),

the De Broglie wavelength works out to 1.2 Å at 100 eV and 1.7 Å at 50 eV. It is therefore possible to move the transition region from 100 eV to lower energies.

There is also a physical reason to do so. Consider the inelastic mean free path of silicon in figure 2.8a. Its minimum is 4 Å at 50 eV for the Penn model, and 3 Å at that energy for the Mermin model. The phonon model assumes that the electron's wave function looks like a plane wave that is several lattice constants long. But because the inelastic scattering mean free path is also on the order of the lattice constant, it is doubtful that this condition is satisfied at 50 eV. It may make more sense to use Mott scattering in the 50–100 eV region. Though Mott scattering also uses the invalid assumption that the electron is a plane wave far from the scattering centre, at least it does not hinge on the assumption that the wave packet is coherent over interatomic distances.

We perform the following sensitivity study. Instead of transitioning from phonon and Mott scattering between 100–200 eV, we choose two different transition regions. Results are given in figure 6.10. We follow the conventional experimental distinction between secondary electrons (SEs) and backscattered electrons (BSEs). In the discussion below, we shall distinguish between the terms “SE” and “secondary electron”. The term SE shall refer to a secondary electron as it is historically known in an experimental context, that is, any electron with less than 50 eV at the time of detection. A secondary electron shall refer to an electron which is liberated from the solid by a primary or a backscattered electron. A secondary electron may have more than 50 eV of energy, though this is uncommon. We could make a similar distinction between BSE and backscattered electrons, but this is not necessary for the purposes of the following discussion. BSEs are all electrons that have more than 50 eV at the moment of detection.

The BSE yield is barely affected by the choice of the low-energy elastic scattering mechanism, as expected. Note, though, that the BSE yield has a “shoulder” at < 200 eV beam energies when the transition region is in the 100–200 eV range. This feature is a direct consequence of our artificial transition between phonon and Mott scattering. It is unlikely to show up in real experiments. The “step” in SE yield at 50 eV is due to the fact that backscattered electrons are classified as SEs when the beam energy is less than 50 eV.

We explain the shape of the SE yields with the following model. SEs escape the material from a thin (few nm) surface layer. At very low beam energies, primary electrons suffer substantial deflection within this surface layer. Coupled to a short inelastic mean free path, this means that every primary electron creates multiple secondary electrons within the surface layer. A higher beam energy means that more secondaries can be created, and the SE yield rises. When




Figure 6.10: (a) secondary electron and (b) backscattered electron yields. The region in which we transition between phonon and Mott scattering is varied.

the beam energy rises further, the inelastic mean free path increases and the deflection from both inelastic and elastic events becomes increasingly forward peaked. The SE yield then reaches a maximum as primary electrons lose most of their energy below the surface layer, in a region from which secondary electrons cannot escape. At high beam energies, the primary and backscattered electrons pierce through the surface layer without being deflected much at all, creating at most one secondary electron on their way through. In this region, we expect the SE yield to be inversely proportional to the inelastic mean free path of the primary electrons, and this is indeed the case.

The fact that the SE yield is slightly higher at high energies for the “Mott only” curve is due to the fact that Mott scattering has a slightly larger transport length than phonon scattering at low energies. The escape layer for SEs is therefore thicker when Mott scattering is used at all energies. This is also why the energy of maximum yield shifts up. To explain why the yield at lower beam energies also goes down, we need to mention the fact that most secondary electrons, after being created, undergo one or more inelastic scattering events before leaving the material. This was already discussed in section 3.3.1, where we saw that the majority of SEs are actually tertiary (or higher order) electrons. Because electrons have a longer range of travel, they can escape the material more easily, but the competing effect is that they can’t create tertiary electrons as effectively. The net result, in this case, is that the SE yield goes down.

6.4 Wave and band structure effects

We have assumed that the electron behaves like a classical point particle, with a well-defined position and momentum, moving through a homogeneous “soup” of material. In reality, the electrons are quantum-mechanical wave packets which are affected by the material’s (crystal) structure. Two physical effects are left out because of this: the band structure and diffraction of the electrons.

6.4.1 Band structure effects

We have assumed a very simple band structure model: electrons are forbidden from having an energy in the band gap, but they behave as free electrons otherwise. In a real solid, there is a complex relationship between the energy and momentum of the electron.

The theory behind Mott scattering uses the vacuum dispersion relation between the electron’s energy and momentum—using the vacuum electron mass and not the effective mass. Similarly, the central quantity in the dielectric function formalism is the electron’s velocity, which is converted to energy using the vacuum dispersion relation. The result is that the mean free path, which we tabulate as function of energy, is “read off” at the incorrect position.

6.4.2 Electron diffraction

In a real crystalline solid, electrons may diffract. This effect is used, for example, in transmission electron microscopy (TEM) or low-energy electron diffraction (LEED) experiments.

It should be noted at this point that electrons, being fermionic, cannot interfere with each other as photons do. An electron in a diffraction experiment must be coherent with itself [157], that is, it is one particle with a very broad wave packet. The electron wave must be broader than the unit cell of the atoms in the sample for diffraction to be significant. Such states can be prepared by keeping the opening angle of the beam small—in other words, by choosing a sufficiently small aperture. This is what happens in TEM and LEED experiments. SEM applications operating at the resolution limit must use small apertures and therefore feature coherent beams too. CD-SEM is such an application, but as it is rarely used with crystalline surfaces, diffraction is unlikely to be a prominent effect. Many other SEM applications use a larger opening angle to increase the current that reaches the sample, thus making the beam less coherent.

For the incident beam, then, it is usually clear whether diffraction can be neglected. What is unclear is what happens after an inelastic scattering event: it is possible that the degree of coherence after an inelastic event can be enough for diffraction to take place. Diffraction after a scattering event would not be immediately obvious in an experimental setting as the electron's direction is random. It may represent an additional elastic scattering mechanism that we have neglected so far.

6.5 Other bulk effects

6.5.1 Bremsstrahlung

Bremsstrahlung refers to the radiation emitted as the electron is decelerated. We might perform a very rough estimate concerning the strength of this effect by considering the stopping power, which is given for silicon in figure 2.8b. The acceleration, a , is related to the stopping power S by $a = -S/m$, where m is the electron mass. We ignore the fact that the electron slows down in discrete events and pretend it happens continuously. The power radiated by Bremsstrahlung can then be computed from the Larmor formula,

$$P = \frac{e^2 a^2}{6\pi\epsilon_0 c^3}. \quad (6.5)$$

Here, e is the electron charge and c is the speed of light. Typical numbers for the stopping power are $S = 50 \text{ eV nm}^{-1}$ at $E = 100 \text{ eV}$. The Larmor formula then states that an additional $5 \times 10^{-7} \text{ eV nm}^{-1}$ are lost to Bremsstrahlung: a negligible amount.

The above is not a very good estimate, as the discrete nature of events means that the instantaneous acceleration is much higher. The radiated power is therefore much higher as well. More accurate numbers can be found in EEDL. The total scattering cross section and stopping power due to ionization (which is present in the simulator) and due to bremsstrahlung are given in figure 6.11. The bremsstrahlung effect is negligible compared to ionization for the energy range of interest.

6.5.2 Defects

In a crystalline material, all kinds of defects (vacancies, interstitials, impurities, dislocations) can form scattering and trapping sites.




Figure 6.11: (a) cross sections and (b) stopping power of bremsstrahlung (solid lines) and ionization (dashed) taken from EEDL for silicon.

In a metal, the amount of scattering due to impurities and dislocations can be measured for Fermi-level electrons as the so-called residual resistance. At room temperature, scattering due to phonons usually dominates. If the temperature is then lowered, the amount of phonon scattering is decreased, and the concentration of vacancies and interstitials decreases. The remaining amount of resistivity is due to impurity and dislocation scattering². As phonon scattering usually dominates near the Fermi energy at relevant temperatures, scattering due to defects is usually ignored in Monte Carlo simulations.

Another effect comes into play in semiconductors. Impurities are classified as electron donors or acceptors depending on their valence relative to the rest of the solid. A similar classification exists for vacancies and interstitials: vacancies usually act as acceptors and interstitials act as electron donors. While a pure semiconductor has no allowed states in the band gap, defects provide additional states which lie in the band gap. When unoccupied, such states form traps which can capture electrons from other energy levels. Defects in semiconductors can also act as scattering centres just like they do in metals.

In addition to point defects discussed above, there are also macroscopic imperfections: dislocations and grain or phase boundaries. These act as additional scattering sites and can also form trap states in the band gap.

Generally speaking, the probability to trap or scatter electrons with defects decreases when their energy increases [74]. For the hot electrons we are dealing

²This discussion disregards superconducting materials.

with here, it is therefore most likely acceptable to ignore these effects.

6.5.3 Polaronic trapping

When an electron moves through a solid, the material polarizes and slows the electron down. The combination of an electron and its polarization cloud is called a polaron [158]. The polaron may form a trapping mechanism which is most effective at low energies.

Ganachaud and Mokrani [58] have proposed an empirical polaron trapping model, with a mean free path given by

$$\lambda^{-1} = \lambda_0^{-1} e^{-\gamma E}. \quad (6.6)$$

E is the energy of the electron, λ_0 and γ are empirical parameters that are fitted to simulation results. This model was also used by Dapor *et al.* [11, 159, 160]. This model arises from no other consideration than that it should be highly effective at low energies and highly ineffective at high energies. It is assumed that the electron loses all kinetic energy in the event of being trapped.

Intuitively, one may expect that polaron trapping is already present in the dielectric function model. The model of Ganachaud and Mokrani behaves very differently than the dielectric function model, though. The mean free path from the dielectric function model tends to infinity for low energies, while the polaron trapping model has a finite mean free path even at zero kinetic energy.

Ganachaud and Mokrani [58] use $\lambda_0 = 1$ nm and $\gamma = 0.25$ eV⁻¹. Dapor *et al.* [11, 159, 160] use similar values for λ_0 but notably lower values for γ , in the range 0.085–0.16 eV⁻¹. In all cases, the trapping mechanism suppresses the secondary electron yield significantly, because most of the secondary electron yield consists of electrons only a few eV above the Fermi level.

6.5.4 Delocalised secondary electron excitation

We currently assume that secondary electrons are created at the exact position where the primary electron lost its energy. In reality, most secondary electrons are created by plasmon decay, which may propagate before decaying.

We estimate the group velocity of a bulk plasmon using equation (2.59),

$$v_g = \frac{\partial \omega}{\partial q} = \frac{3v_f^2}{\omega_p} q. \quad (6.7)$$

v_f and ω_p are the material's Fermi velocity and plasmon energy, respectively. The maximal value of q is given by the point where the plasmon dispersion line

intersects the region of single excitation in figure 2.7. Plasmons with higher q decay almost immediately due to Landau damping. For a free-electron gas with $\hbar\omega_p = 17$ eV, the maximal value of q is 11 nm^{-1} and the corresponding group velocity is $6.2 \times 10^6 \text{ m s}^{-1}$.

The lifetime of a plasmon can be estimated from the width of the plasmon peak. Abril *et al.* [48] use $\gamma = 4.245$ eV for the width of the plasmon peak in silicon, which leads to a lifetime of $\hbar/\gamma \approx 1.5 \times 10^{-16}$ s. Combined with the group velocity found above, we find that plasmons in silicon travel a maximum of 0.9 nm. This is a significant number, though it must be emphasised that it is an upper bound. Plasmon excitation with lower values of q is much more likely, thanks to the fact that the energy-loss function $\text{Im}[-1/\epsilon]$ decreases for large q and the factor $1/q$ in equation (6.1).

6.6 Interface effects

6.6.1 Shape of the work function potential

We have assumed a simple step function potential at the material boundary (section 2.3). This led to a quantum-mechanical transmission probability, even when the electron has enough energy to escape the material. It is likely that the actual potential seen by the electron is smoother than a step function. We may expect that a smoother potential will lead to a larger transmission coefficient.

For example, Villarrubia *et al.* [17] replace the step function potential by a logistic function,

$$U(x) = \frac{\Delta U}{1 + e^{-x/w}}, \quad (6.8)$$

where ΔU is the height of the barrier and w is the width. This leads to the following reflection coefficient:

$$R = \frac{\sinh^2(\pi w(k_1 - k_2))}{\sinh^2(\pi w(k_1 + k_2))}. \quad (6.9)$$

$k_1 = \sqrt{2mE \cos^2 \theta}/\hbar$ and $k_2 = \sqrt{2m(E \cos^2 \theta) - \Delta U}/\hbar$ are the electron wave numbers before and after the barrier; θ is the angle of the electron with respect to the surface normal before it reaches the boundary. When $\Delta U > E \cos^2 \theta$, only reflection is possible. If the characteristic width w is small, this reflection coefficient reduces to the step potential case. In the limit that w tends to infinity, this becomes a classical barrier.

Unfortunately, it is not a priori clear what the width w of the barrier should be. It is most likely on the order of the size of a single atom, that is, about 1 Å. We may get a feel for the importance of this barrier by considering the extreme case in which $w \rightarrow \infty$, that is, by using a classical barrier. The resulting yields (for silicon) can be seen in figure 6.12. The SE yield is increased by 30% for all incident energies, the BSE yield is unaffected. A linescan over a 40 nm wide and 40 nm high silicon line on a silicon substrate can be seen in figure 6.13. The shape of the linescan is barely modified. The width of the edge blooming effect, measured in the same way as in section 3.6.2, becomes slightly larger. It is 2.73 nm in case of a sharp quantum mechanical step and becomes 2.93 nm when a classical barrier is used.

6.6.2 Alignment of Fermi energies

When two materials are placed into contact, their Fermi levels align. The electronic structure locally changes and band bending occurs.

For simplicity, we will now consider two metals placed into contact. There are two effects which are ignored in our simulations. They are both illustrated in figure 6.14.

First, the potential barrier between materials is altered. The current models take the vacuum level as a constant. When an electron enters or leaves a material, it crosses a potential barrier $U = F + \Phi$, where F is the material's Fermi energy and Φ is the work function. When it moves from material 1 to material 2, because its energy with respect to the vacuum level must not change, the electron crosses a barrier equal to $\Delta U = U_2 - U_1 = F_2 - F_1 + \Phi_2 - \Phi_1$. However, if the Fermi levels are aligned, this barrier should be equal to $F_2 - F_1$.

The second effect is closely related. In order to maintain energy balance, there must be an electric field outside the material when these two materials are placed into contact. This electric field accounts for the difference in work functions between the materials. Electric fields are currently not implemented in the simulator. If they are implemented in the future, it would be best to change the boundary crossing mechanism to reflect the situation of figure 6.14b.

6.6.3 Surface states

In an ideal, infinite crystal, the energy eigenstates of electrons are Bloch waves. If the crystal has a surface, two kinds of states are possible. The first type is similar to a Bloch wave inside the crystal, but exponentially decays on the outside. The second type is a surface state, localised at the edge of the crystal. Passing




Figure 6.12: (a) secondary electron and (b) backscattered electron yields of silicon. We use a quantum mechanical step barrier and a classical (infinitely smooth) barrier at the interface to obtain the transmission coefficients. The barrier energy $U = 12.48$ eV (see appendix A).

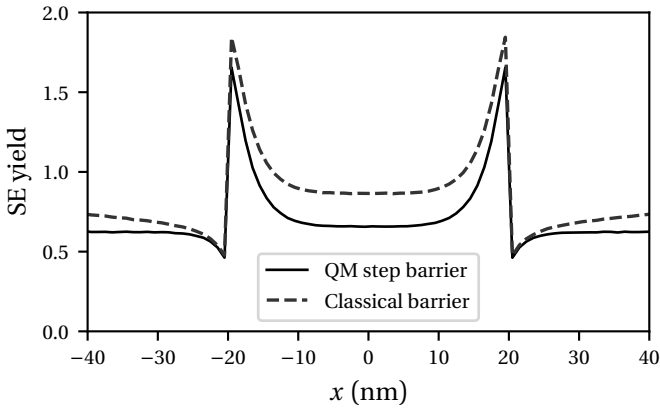


Figure 6.13: Linescan of an isolated 40 nm high and 40 nm wide silicon line on a silicon substrate. The feature is centered around $x = 0$. The electron beam is infinitely sharp and has a 500 eV landing energy. The intensity on the vertical axis is derived from the emitted secondary electrons. We use a quantum mechanical step barrier and a classical (infinitely smooth) barrier at the interface to compute the transmission coefficients.




Figure 6.14: Relative positioning of bands. Two metals are placed in contact with each other, the vertical axis represents the energy. $F_{1,2}$ and $\Phi_{1,2}$ represent the Fermi energy with respect to the bottom of the band and the work function, respectively, for each material. (a) represents the current simulator, in which the vacuum levels are aligned. (b) represents the situation when the Fermi levels are aligned. ΔU is the jump in inner potential that an electron experiences when moving between two materials. In situation (b), an electric field with potential $\Delta\Phi$ is set up outside the materials.

electrons may be trapped in such surface states, which we have ignored in this thesis.

Another, more practical issue is that clean crystal surfaces are usually chemically unstable. Most metals and semiconductors form oxidation layers on their surface when exposed to air. Because surface traps due to dangling bonds can severely reduce device performance, silicon wafers are often terminated with hydrogen in the semiconductor industry. Surface termination can significantly influence the work function and the amount of band bending [161, 162] and can therefore have a substantial impact on the emission of secondary electrons. There are two related issues here: is it possible to model how the surface influences electron scattering, and are experimental conditions sufficiently controlled that the surface really is what we expect it to be?

6.7 Electron–electron interaction

The electrons in the simulation are assumed to be independent from each other. In reality, their charges will repel each other, leading to, amongst others, charging effects. The theoretical work of Cazaux [163–165] must be mentioned in this




Figure 6.15: The time it takes for an electron to slow down to 50 eV, as function of its initial energy. The material is silicon, the model used is Penn’s algorithm.

regard. Several authors have integrated charging into a Monte Carlo framework [17, 166–169]. Here, we will limit ourselves to a brief description of the physics involved, ignoring the technicalities of the implementation.

6.7.1 Interaction with residual charge from previous illumination

When an electron lands on a sample, it may be influenced by the charge distribution left behind by previous electron cascades.

In most practical situations, we may assume that the previous cascade has finished scattering long before the next electron in the beam reaches the sample. Figure 6.15 shows the time it takes for an electron to slow down from a given starting energy to 50 eV. This was computed from the stopping power in silicon (figure 2.8b). The stopping time is computed as

$$t = \int_{50\text{eV}}^{E_0} dE \sqrt{\frac{m}{2E}} \frac{1}{S(E)}, \quad (6.10)$$

where m is the electron mass, $S(E)$ is the stopping power as function of energy and E_0 is the starting energy. An electron cascade usually ends in 1×10^{-13} s. The highest beam currents seen in most SEM applications are on the order of 10 nA, which corresponds to one electron every 1×10^{-11} s on average.

If the sample is conductive, the residual charge can be screened out over short distances, it is easy for the charges to recombine and any charge imbalance (if the electron yield is not equal to one) is easily replenished. If the sample is nonconductive, the residual charge from all previous cascades cannot be

screened effectively, and the Monte Carlo cascade is affected by the resulting electric field. This diffusion timescale is much longer than the timescale of the Monte Carlo cascade [170], so the residual charge may be considered static for the duration of a single electron cascade.

6.7.2 Interaction with secondary electrons from the same cascade

While electrons from successive cascades are not simultaneously active, the same cannot necessarily be said of electrons that are part of the same cascade. If a secondary electron is created by direct impact from a primary electron, these electrons may briefly be very close together. This does two things: first, it introduces an additional repulsive force between the two; second, the existing scattering mechanisms (such as dielectric function scattering) may behave differently.

The distance that two electrons need to be apart for them not to influence each other can be estimated by the Thomas-Fermi screening length. It is given by

$$L_{\text{TF}}^{-2} = k_{\text{TF}}^2 = \frac{4}{a_0} \left(\frac{3n}{\pi} \right)^{1/3}. \quad (6.11)$$

In silicon, $n \approx 2 \times 10^{23} \text{ cm}^{-3}$ and $L_{\text{TF}} \approx 0.5 \text{ \AA}$.

Most secondary electrons are not created by direct ionization but by plasmon decay, which gives the primary electron time to move away. As discussed in section 6.5.4, the lifetime of a plasmon in silicon is on the order of $1.5 \times 10^{-16} \text{ s}$. An electron with 100 eV travels almost 6 \AA in that time, much more than the screening length. Metals usually have sharper plasmon peaks, corresponding to a longer lifetime. The plasmon will also move before it decays (see section 6.5.4), most likely in a different direction than the primary electron.

It is also possible for electrons to get close to each other randomly, after several scattering events, but the probability of this happening is very low.

6.7.3 Dynamics of residual charge

Residual charge from previous cascades may diffuse on a much slower time scale. Raftari *et al.* [170] have developed a very sophisticated drift-diffusion-reaction model which attempts to describe most of the electron cascade as a diffusion process. The Monte Carlo method should be more accurate at high energies, but such a diffusion model could very well describe the slow time-scale dynamics of residual charges without enough energy to escape the sample.

If the electric field becomes very strong, dielectric breakdown is also possible. This phenomenon has been well studied from a theoretical perspective [83]. It has also been implemented in Monte Carlo models [167], under the assumption that the breakdown represents an instantaneous rearrangement of charge that does not influence the “microscope electrons”. It is unclear how accurate the assumption is, as dielectric breakdown in devices can last several nanoseconds [171].

Dielectric breakdown may also damage the material, locally changing its optical properties and therefore changing the inelastic scattering of electrons. This effect appears to have been ignored in Monte Carlo simulations so far.

6.8 Experimental issues

The previous sections discuss difficulties that are inherent in the theory. There are also problems on the experimental side. One problem lies with the input data for the model: parameters such as the optical dielectric function or the speed of sound are obtained experimentally, each with their own inherent uncertainties. Additional problems arise due to the fact that we are trying to mimic a real electron microscope, but a real electron microscopy application is never perfect. We discuss these issues in this section.

6.8.1 Native oxides

Many materials of interest oxidize when exposed to air. This native oxide is rarely cleaned off before inspection with an SEM, which means there is usually an oxide layer with an unknown thickness on the sample.

We estimate the contribution of native oxides by simulating electron yields. These are shown in figure 6.16: the substrate is silicon and the oxide layer is SiO_2 . SEs are all detected electrons with an energy less than 50 eV; BSEs are all others. The SE yield, which is known to be sensitive to the surface condition, is substantially altered by a layer of SiO_2 . The absolute value is increased considerably and the energy of maximal yield shifts upwards. The difference in BSE yields is less significant. A more detailed analysis (not shown here) indicates that the difference in BSE yields is mostly due to high-energy secondary electrons with energies in the 50–200 eV range.




Figure 6.16: (a) secondary electron and (b) backscattered electron yields. Solid curves are clean silicon; the dashed and dash-dotted curves have a thin layer of SiO_2 added on top.

6.8.2 Surface contamination

An electron microscope does not have a perfect vacuum. In fact, most SEMs only achieve a high vacuum of around 1×10^{-6} – 1×10^{-7} mbar. Hydrocarbon contamination from the chamber may adsorb to the surface of the sample. The sample itself may also be a source of hydrocarbons, which it may have acquired during preparation, storage or transfer. These hydrocarbons may be dissociated by electron irradiation—either from the primary electron beam or by secondary electrons escaping the sample—and stick to the surface.

As surface contamination often mostly consists of carbon, we perform a similar study as for native oxides, using graphite as the surface layer. The results can be found in figure 6.17. In these simulations, the graphite affects the secondary electron yield much less significantly than the oxide layer. While SiO_2 seems to enhance the SE yield, graphite appears to suppress it a little bit.

6.8.3 Sample damage

Many interesting samples can be damaged by electron irradiation, due to several effects [172].

The first effect we discuss is direct damage to the sample due to inelastic scattering. This is primarily a problem in organic compounds: molecules are excited or ionized by an inelastic event and chemical bonds are broken. The molecule may then change shape or diffuse out of the material. A related phe-




Figure 6.17: (a) secondary electron and (b) backscattered electron yields. Solid curves are clean silicon; the dashed and dash-dotted curves have a thin layer of graphite added on top.

nomenon is desorption induced by electronic transitions [173], in which repulsive electronic states are directly induced in surface molecules. Photoresists, for example, are known to shrink substantially when exposed to an electron beam [174, 175]. Verduin [12, figure 2.19] has shown that this effect is significant at very low electron doses on the order of 1 electron per nm^2 .

A second effect is sample heating. A small amount of the energy from the incoming electron beam leaves the sample in the form of escaping electrons. Most of this is carried away by backscattered electrons, which usually consist of 20–50% of the incoming beam, depending on the sample material. The rest of the energy is deposited in the sample as heat. In scanning electron microscopy applications, this effect is a few kelvin at most if the beam is stationary [95] and less if the beam scans. Thin samples, in which the heat can only dissipate in two dimensions, may suffer from more heating.

When an electron scatters elastically, for example due to Mott scattering in our simulator, it is deflected. Due to momentum conservation, a small amount of recoil energy must be transferred to the atom with which the scattering took place. This is usually a very small amount, but in some cases it can lead to the atom being displaced. The energy required for displacing an atom is smallest on the surface. In practice, this effect is only seen in transmission electron microscopy, where primary electrons have very high (1×10^2 keV) energy. As most momentum transfer takes place in the direction of the incident electron, the back side of a sample in a TEM may suffer from this effect, where it is known as

sputtering.

In experimental conditions, damage is sometimes reduced by applying surface coatings to act as a barrier preventing the escape of volatile elements. Conductive surface coatings are also used to mitigate the effects of charging. Samples in transmission electron microscopes are sometimes cooled to liquid nitrogen temperatures. This is also found to reduce damage, most likely due to the reduced mobility of atoms.

6.9 Discussion

The previous sections discuss inherent model uncertainties, effects that have been modelled in literature but represent only small corrections, and effects that we do not know how to model within the context of our Monte Carlo simulator. We have made an effort to quantify the importance of these effects where possible. The question now arises which topics should be addressed first in a follow-up study.

For the purposes of scanning electron microscopy, it is probably unrealistic to expect an exact match between simulation and experiment. Figures 6.16 and 6.17 show the uncertainties in electron yields due the presence of oxides and carbon contamination. These factors are not controlled in experimental situations, and the uncertainties due to these exceed the inherent model uncertainties, see figure 3.10, considerably. That does not mean it is unnecessary to improve the modelling side. In order to model charging effects correctly, it is important to get the charge balance right. This is not just the yield, which represents the overall charge of the sample, but also the positions of all charges which set up the electric field. The range that secondary electrons can travel could be an important factor here.

The most significant model uncertainties remain in the < 100 eV energy region. The most low-hanging fruit we have seen is the incorporation of exchange effects, which can cause up to a factor 2 discrepancy in the inelastic mean free path. We have seen two approximations in literature. At present, there is not much reason to prefer either one of them over the other, but they give broadly similar results anyway. It could also be very interesting to give higher-order Born approximations a closer look. The second-order approximation produces unphysical results at low energies in our example. Higher orders should be taken into account to correct for these. Before the second order starts to break down at approximately 60 eV, we found a 20% deviation in mean free path compared to the first order. This deviation is only expected to increase towards lower

energies. The third topic to consider is elastic scattering at low energies. We have used a transition between phonon and Mott scattering in the 100–200 eV range, but it is possible that this should be altered. One way to calibrate this would be to measure the range that secondary electrons can travel, possibly by measuring the size of the edge blooming effect.

If charging effects are of specific interest, electrons should be tracked until they reach the bottom of the conduction band (or the Fermi level in case of metals). In that case, trapping by defects or polarons may have a substantial influence as well. Such effects are most effective at very low energies. For modelling charging effects, it is not sufficient to implement only an electric field solver. The dynamics (*e.g.* diffusion) of the residual charge change the field over time. It is likely acceptable to treat the Monte Carlo cascade as “instantaneous” compared to the diffusion time scale, though.

Several other effects that we have discussed, such as the tensorial nature of the dielectric function, a direction-dependent band structure, and diffraction effects should be of lower priority. It would take a lot of effort to incorporate these into the present framework, and their effects are likely to be averaged out as an electron scatters through the material. Though diffraction can be observed experimentally, in practice this only happens when both the sample and experimental conditions are optimized for this purpose, such as in low-energy electron diffraction (LEED) experiments.

With all these uncertainties in mind, is it worth doing simulations in the first place? The answer is yes. The results of chapters 4 and 5 are not sensitive to specific model parameters. The same goes for most of the illustrative results in chapter 3. These results have meaning of their own and we can truthfully say we have learned something. As Richard Hamming once famously said, the purpose of computing is insight, not numbers [176]. With the current models, our simulator can provide us with plenty of valuable insight and understanding.

Chapter 7

Conclusion

The goal of this thesis was to develop a Monte Carlo simulator of electron-matter interaction with accurate physical models. This simulator was intended for use in scanning electron microscopy (SEM) applications, but it can also be used for related purposes.

We carefully reviewed the physical models in literature before creating an open-source simulation package that can run both on graphics processing units (for speed) and central processing units (for flexibility). We chose a hybrid set of models comprising the Penn optical data model, Mott scattering, and electron-acoustic phonon scattering. The first describes how electrons lose energy. This energy loss was used as input for a model for how secondary electrons are excited. The second and third models describe (quasi-)elastic scattering, and we used an (arbitrary) 100–200 eV transition region between the two. Alternative options for this transition region were explored in the final chapter. The way in which we deal with low-energy elastic scattering leaves some things to be desired. In addition to the aforementioned transition region, the phonon model depends on parameters (such as the deformation potential or lattice constant) that are not easy to measure or do not exist for many materials. We used the phonon model because we believe it makes more physical sense than many alternatives in literature, such as using Mott scattering all the way down to 0 eV.

The simulator allowed us to gain a basic understanding of the interaction volume. The “secondary electron” peak in the energy spectrum is actually dominated by the presence of tertiary (and higher order) electrons. In addition, the interaction volume only has a “pear shape” when the beam landing energy is at least several keV. We also compared the simulator to experimental electron yields and found a fairly good agreement.

With a good simulator in place, we investigated an effect that is often ignored in the literature of SEM simulations: the contribution of surface plasmons to secondary electron emission. The inclusion of surface plasmons allows for non-local secondary electron excitation: a primary electron outside the material may excite a surface plasmon which can decay to an electron-hole pair. We found that the inclusion of surface plasmons is important if one is interested in energy spectra, but not in the *number* of emitted electrons. More importantly, the shape of a low-voltage SEM signal is barely affected. The SEM signal is only significantly affected when a primary electron is very close (sub-ångström distance) to the surface and travels nearly parallel to it. This situation is so rare in practice that it should have no appreciable influence on low-voltage CD-SEM images.

Having established that, it was time to apply the simulator to a more practical situation in lithography: the measurement of line-edge roughness. If we consider an isolated line, it turns out that the edge position observed by the SEM is given by the outermost extrusion of the rough sidewall along the vertical direction. We called this the “projection model”. Due to this projection effect, the measured line-edge roughness is less than the roughness of any particular slice of the sidewall. A lower line-edge roughness is measured for thicker lines, even if the actual sidewall roughness is the same. For a dense line-space pattern, the SEM is less sensitive to the lower layers, and the edge position detected by the SEM lies between the position of the top slice of the sidewall and the projection model. Lines are “isolated enough” for the projection model to hold when the trench between the lines is larger than the height of the lines themselves.

In the final chapter, we analysed the physical effects that are still missing from the simulator as it currently stands. There is plenty of room for improvement on the modelling side. As we point out in that chapter, though, there are more fundamental issues when modelling SEM images in particular. One usually does not know the true composition of samples of interest. Many metallic samples have native oxides of unknown thickness, samples in all SEMs suffer from contamination due to hydrocarbon deposition from the imperfect vacuum conditions, and many samples are susceptible to damage. Simulators like the one developed in this thesis are only able to tackle the forward problem: they can only give us a SEM image if the sample is known exactly upfront. The inverse problem, which is the one we would like to solve, is much more difficult and perhaps even impossible for SEM.

Acknowledgements

This thesis would not have existed without the help of many people, to whom I would like to express my deepest gratitude.

First of all, I would like to thank my promotor Kees, who gave me the opportunity to do a PhD project in what was (at the time) called the charged particle optics group. Thank you for your constant enthusiasm and support through these four years. I have enjoyed our many long discussions, not just about my work but also the broader context of microfabrication and the rest of the world. I am also very grateful to my second promotor Pieter. You were always able to provide a fresh perspective, and I was always deeply impressed by your very deep insight into seemingly anything related to physics.

This project would not have taken place without the support from ASML. In particular, I would like to thank Wim, Sander and Coen for giving me the opportunity to start this project. Special thanks to Sander for keeping a keen interest in my progress and asking critical questions about the broader context. I also want to thank many team members at ASML, who, despite certain IP restrictions, were always interested in my work and happy to contribute wherever possible. Thomas, thank you for your enthusiasm, encouragement, helping me kickstart chapter 5 and providing creative and out-of-the-box ideas throughout that project. Thanks to Benoit, Ton, Shakeeb, Willem, Jasper, Ruben and Bernardo for many interesting discussions, and also for using the software and providing valuable feedback.

It has been a pleasure to do my PhD project in the MInT (formerly CPO) group at TU Delft. I had many interesting discussions with Annemarie and Kerim, who were working on similar simulation projects as me. Thank you Ryan and Aya, my office mates, for a very nice and friendly environment even though I was out two days per week. In particular, I want to thank Ryan (again), Yoram and Mathijs for many extended coffee breaks while I had simulations running. It has been great hanging out with you, also outside of work and in times of

COVID-19 (online or otherwise) and after my contract ended. And to all other people in the group, including Maurice, Wilco, Aditi, Sangeetha, Gaudhaman, Lixin, Gerward, Shammi, Yan, Laura, Xiaoli, Neda, Xin, Vidya, Marco, Nicolò, Anjella, Jacob, Daan, Elizabeth, Carel, Han, Johan, Ruud, Dustin and many others: thank you for warmly welcoming me and making my time there unforgettable.

I am deeply indebted to Thomas Verduin, who left the group just before I joined but was kind enough to help me during my master project. He taught me how GPU programming works, and his GPU simulator was the starting point for this thesis.

A PhD project is not complete without side projects, which is why I want to thank Evgenia, Alexey and Vina for involving me in a very interesting one. I really appreciated your encouragements that my contributions were useful despite some inevitable accuracy issues. It was great seeing the related project of Sonia and Olivier, which inspired me to keep going. In particular, I want to thank Marijke for taking the initiative to do experiments with me. I enjoyed our time together owing to your infectious excitement and our many chocolate muffins. I hope we will collaborate again in the future!

Milo, Bjarni and Ruud, it has been great sharing this PhD journey with you even though we were all at different universities. I have enjoyed our trips together and hope we will do many more. Our heroic descent down the mountain will never be forgotten!

Last but not least, I would like to thank my parents and sister. Wim, Angeline and Nicole, thank you for your constant support and encouragement through these many years. I could not have done it without you.

Appendix A

Material parameters

The material parameters used in this thesis are given in table A.1. The names of most materials speak for themselves. PMMA stands for poly(methyl methacrylate), with chemical formula $C_5O_2H_8$. C stands for graphite. The symbols represent the following physical quantities and units:

- ρ_m : density (g cm^{-3})
- V : width of the valence band (eV)
- G : band gap (eV)
- χ : electron affinity (eV)
- F : Fermi energy with respect to the bottom of the band (eV)
- Φ : work function (eV)
- a : lattice constant (\AA)
- c_L : longitudinal speed of sound (km s^{-1})
- c_T : transverse speed of sound (km s^{-1})
- Ξ_L : Longitudinal deformation potential (eV)
- Ξ_T : Transverse deformation potential (eV)
- Ξ : Isotropic deformation potential (eV)
- ρ : Resistivity ($1 \times 10^{-8} \Omega \text{ m}$)

The Fermi energy of graphite was estimated using a method proposed by Kieft and Bosch [10]. They estimate it as $F = I - \Phi$, where I is the ionization potential and Φ is the work function.

Table A.1: Material parameters used for the simulations in this thesis.

	Si	Al	Cu	PMMA	SiO ₂	C
ρ_m	2.329 [177]	2.70 [177]	8.96 [177]	1.2 [177]	2.648 [177]	2.2 [177]
V	7.27 [178]			0 ^a	9 [179]	
G	1.11 [180]			5.6 [181]	8.9 [179]	
χ	4.1 [182]			2.5 [10]	0.9 [183]	8.64 ^b
F		11.63 [180]	7.00 [180]			
Φ		4.08 [184]	4.7 [184]			4.6 [185]
a	5.43 [186]	4.05 [187]	3.62 [188]	1.54 ^c	5.45 [79]	2.46 [189]
c_L	9.13 [190]	6.42 [177]	4.76 [177]	2.75 [191]	5.97 [192]	3.29 [193]
c_T	5.84 [190]	3.04 [177]	2.33 [177]	1.39 [191]	3.76 [192]	1.80 [193]
Ξ_L	9.2 [194]					
Ξ_T	5.0 [194]					
ρ		2.82 [184]	1.7 [184]	2 ^d	2.6 [79]	16 [195]

^a Value unknown, zero used by default.^b Estimated from ionization potential 13.24 eV [196].^c C-C bond length used.^d Unknown, tuned such that reasonable mean free paths were obtained.

Curriculum Vitæ

L. C. P. M. van Kessel

26-05-1993 Born in Veldhoven, The Netherlands.

Education

2005–2011 Gymnasium
Sondervick College, Veldhoven

2011–2014 Bachelor in Physics & Astronomy (cum laude)
Radboud University Nijmegen

2014–2017 Master in Physics & Astronomy (cum laude)
Radboud University Nijmegen

Awards

2020 Karel Urbánek best student paper award

List of publications

- L. van Kessel, C.W. Hagen, and P. Kruit, “Surface effects in simulations of scanning electron microscopy images”, Proceedings of SPIE **10959**, 109590V (2019).
- L. van Kessel, C.W. Hagen, and P. Kruit, “Surface effects in simulations of scanning electron microscopy images”, Journal of Micro/Nanolithography, MEMS, and MOEMS **18**, 044002 (2019).

- L. van Kessel, T. Huisman, and C. W. Hagen, “Understanding the influence of 3D sidewall roughness on observed line-edge roughness in scanning electron microscopy images”, *Proceedings of SPIE* **11325**, 113250Z (2020).
- L. van Kessel, T. Huisman, and C. W. Hagen, “Understanding the influence of three-dimensional sidewall roughness on observed line-edge roughness in scanning electron microscopy images”, *Journal of Micro/Nanolithography, MEMS, and MOEMS* **19**, 034002 (2020).
- L. van Kessel and C. W. Hagen, “Nebula: Monte Carlo simulator of electron-matter interaction”, *SoftwareX* **12**, 100605 (2020).
- J. Löffler, J. Thomet, M. Belhaj, L. van Kessel, C. W. Hagen, C. Ballif, and N. Wyrsh, “Monte Carlo modeling of electron multiplication in amorphous silicon based microchannel plates”, in *2019 IEEE nuclear science symposium and medical imaging conference (NSS/MIC)* (2019), pp. 1–6.
- O. Lugier, A. Troglia, N. Sadegh, L. van Kessel, R. Bliem, N. Mahne, S. Nannarone, and S. Castellanos, “Extreme ultraviolet photoelectron spectroscopy on fluorinated monolayers: towards nanolithography on monolayers”, *Journal of Photopolymer Science and Technology* **33**, 229–234 (2020).
- R. Lane, Y. Vos, A. H. G. Wolters, L. van Kessel, S. Elisa Chen, N. Liv, J. Klumperman, B. N. G. Giepmans, and J. P. Hoogenboom, “Optimization of negative stage bias potential for faster imaging in large-scale electron microscopy”, *Journal of Structural Biology: X* **5**, 100046 (2021).

Bibliography

- [1] J. C. Maxwell, *A treatise on electricity and magnetism* (Oxford: Clarendon Press, 1873).
- [2] J. J. Thomson, “Cathode rays”, *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* **44**, 293–316 (1897).
- [3] P. W. Hawkes, “Ernst Ruska”, *Physics Today* **43**, 84 (1990).
- [4] G. E. Moore, “Cramming more components onto integrated circuits”, *Electronics* **38**, 114–117 (1965).
- [5] L. Zheng-Ming and A. Brahme, “An overview of the transport theory of charged particles”, *Radiation Physics and Chemistry* **41**, 673–703 (1993).
- [6] Lord Kelvin, “Nineteenth century clouds over the dynamical theory of heat and light”, *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* **2**, 1–40 (1901).
- [7] N. Metropolis, “The beginning of the Monte Carlo method”, *Los Alamos Science* **15**, LA-UR-88-9067, 125–130 (1987).
- [8] D. Hebbard and P. Wilson, “The effect of multiple scattering on electron energy loss distributions”, *Australian Journal of Physics* **8**, 90 (1955).
- [9] Y. G. Li, S. F. Mao, H. M. Li, S. M. Xiao, and Z. J. Ding, “Monte Carlo simulation study of scanning electron microscopy images of rough surfaces”, *Journal of Applied Physics* **104**, 064901 (2008).
- [10] E. Kieft and E. Bosch, “Refinement of Monte Carlo simulations of electron–specimen interaction in low-voltage SEM”, *Journal of Physics D: Applied Physics* **41**, 215310 (2008).

- [11] M. Dapor, M. Ciappa, and W. Fichtner, “Monte Carlo modeling in the low-energy domain of the secondary electron emission of polymethylmethacrylate for critical-dimension scanning electron microscopy”, *Journal of Micro/Nanolithography, MEMS, and MOEMS* **9**, 023001 (2010).
- [12] T. Verduin, “Quantum noise effects in e-beam lithography and metrology”, PhD thesis (Technische Universiteit Delft, 2017).
- [13] A. M. M. G. Theulings, “Optimisation of photon detector tynode membranes using electron-matter scattering simulations”, PhD thesis (Technische Universiteit Delft, 2020).
- [14] S. Babin, S. Borisov, A. Ivanchikov, and I. Ruzavin, “CHARIOT: software tool for modeling SEM signal and e-beam lithography”, *Physics Procedia* **1**, 305–313 (2008).
- [15] H. Demers, N. Poirier-Demers, A. R. Couture, D. Joly, M. Guilmain, N. De Jonge, and D. Drouin, “Three-dimensional electron microscopy simulation with the CASINO Monte Carlo software”, *Scanning* **33**, 135–146 (2011).
- [16] *PENELOPE-2014, a code system for Monte Carlo simulation of electron and photon transport*, NEA/NSC/DOC(2015)3 (2015).
- [17] J. S. Villarrubia, A. E. Vladár, B. Ming, R. J. Kline, D. F. Sunday, J. S. Chawla, and S. List, “Scanning electron microscope measurement of width and shape of 10nm patterned lines using a JMONSEL-modeled library”, *Ultramicroscopy* **154**, 15–28 (2015).
- [18] J. Löffler, J. Thomet, M. Belhaj, L. van Kessel, C. W. Hagen, C. Ballif, and N. Wyrsh, “Monte Carlo modeling of electron multiplication in amorphous silicon based microchannel plates”, in 2019 IEEE nuclear science symposium and medical imaging conference (NSS/MIC) (2019), pp. 1–6.
- [19] R. Eisberg and R. Resnick, *Quantum physics of atoms, molecules, solids, nuclei, and particles*, Second edition (Wiley, 1985).
- [20] G. Grosso and G. P. Parravicini, *Solid state physics* (Academic Press, 2000).
- [21] R. H. Ritchie, “Interaction of charged particles with a degenerate Fermi-Dirac electron gas”, *Physical Review* **114**, 644–654 (1959).

- [22] W. W. M. Allison and J. H. Cobb, “Relativistic charged particle identification by energy loss”, *Annual Review of Nuclear and Particle Science* **30**, 253–298 (1980).
- [23] C. J. Tung, Y. F. Chen, C. M. Kwei, and T. L. Chou, “Differential cross sections for plasmon excitations and reflected electron-energy-loss spectra”, *Physical Review B* **49**, 16684 (1994).
- [24] A. Vaglio Pret, M. Kocsis, D. De Simone, G. Vandenberghe, J. Stowers, A. Giglia, P. de Schepper, A. Mani, and J. J. Biafore, “Characterizing and modeling electrical response to light for metal-based EUV photoresists”, *Proceedings of SPIE* **9779**, 977906 (2016).
- [25] S. Bhattarai, A. R. Neureuther, and P. P. Naulleau, “Study of energy delivery and mean free path of low energy electrons in EUV resists”, *Proceedings of SPIE* **9779**, 97790B (2016).
- [26] H. Shinotsuka, S. Tanuma, C. J. Powell, and D. R. Penn, “Calculations of electron stopping powers for 41 elemental solids over the 50 eV to 30 keV range with the full Penn algorithm”, *Nuclear Instruments and Methods in Physics Research B* **270**, 75–92 (2012).
- [27] H. T. Nguyen-Truong, “Low-energy electron inelastic mean free path in materials”, *Applied Physics Letters* **108**, 172901 (2016).
- [28] H. T. Nguyen-Truong, “Electron inelastic mean free path at energies below 100 eV”, *Journal of Physics Condensed Matter* **29**, 215501 (2017).
- [29] E. D. Palik, *Handbook of optical constants of solids, five-volume set* (Elsevier, 1997).
- [30] B. Henke, E. Gullikson, and J. Davis, “X-ray interactions: photoabsorption, scattering, transmission, and reflection at $E = 50\text{--}30,000$ eV, $Z = 1\text{--}92$ ”, *Atomic Data and Nuclear Data Tables* **54**, 181–342 (1993).
- [31] W. S. Werner, K. Glantschnig, and C. Ambrosch-Draxl, “Optical constants and inelastic electron-scattering data for 17 elemental metals”, *Journal of Physical and Chemical Reference Data* **38**, 1013–1092 (2009).
- [32] M. N. Polyanskiy, *Refractive index database*, <https://refractiveindex.info> (visited on 2020-06-18).
- [33] R. F. Egerton, “Electron energy-loss spectroscopy in the TEM”, *Reports on Progress in Physics* **72**, 016502 (2008).

- [34] D. Chandler-Horowitz and P.M. Amirtharaj, “High-accuracy, midinfrared ($450\text{ cm}^{-1} \leq \omega \leq 4000\text{ cm}^{-1}$) refractive index values of silicon”, *Journal of Applied physics* **97**, 123526 (2005).
- [35] C. Schinke, P. Christian Peest, J. Schmidt, R. Brendel, K. Bothe, M. R. Vogt, I. Kröger, S. Winter, A. Schirmacher, S. Lim, H. T. Nguyen, and D. MacDonald, “Uncertainty analysis for the coefficient of band-to-band absorption of crystalline silicon”, *AIP Advances* **5**, 067168 (2015).
- [36] J. Lindhard, “On the properties of a gas of charged particles”, *Det Kongelige Danske Videnskabernes Selskab, Matematisk-fysiske Meddelelser* **28** (1954).
- [37] J. Sólyom, *Fundamentals of the physics of solids: normal, broken-symmetry, and correlated systems*, Vol. 3 (Springer, 2010).
- [38] K. L. Kliewer and R. Fuchs, “Lindhard dielectric functions with a finite electron lifetime”, *Physical Review* **181**, 552–558 (1969).
- [39] N. D. Mermin, “Lindhard dielectric function in the relaxation-time approximation”, *Physical Review B* **1**, 2362–2363 (1970).
- [40] A. Howie and R. M. Stern, “The optical potential in electron diffraction”, *Zeitschrift fur Naturforschung - Section A Journal of Physical Sciences* **27**, 382–389 (1972).
- [41] R. H. Ritchie and A. Howie, “Electron excitation and the optical potential in electron microscopy”, *Philosophical Magazine* **36**, 463–481 (1977).
- [42] D. R. Penn, “Electron mean-free-path calculations using a model dielectric function”, *Physical Review B* **35**, 482–486 (1987).
- [43] J. C. Ashley, “Interaction of low-energy electrons with condensed matter: stopping powers and inelastic mean free paths from optical data”, *Journal of Electron Spectroscopy and Related Phenomena* **46**, 199–214 (1988).
- [44] J. C. Ashley, “Energy loss rate and inelastic mean free path of low-energy electrons and positrons in condensed matter”, *Journal of Electron Spectroscopy and Related Phenomena* **50**, 323–334 (1990).
- [45] Z. J. Ding and R. Shimizu, “Inelastic collisions of kV electrons in solids”, *Surface Science* **222**, 313–331 (1989).
- [46] D. J. Planes, R. Garcia-Molina, I. Abril, and N. R. Arista, “Wavenumber dependence of the energy loss function of graphite and aluminium”, *Journal of Electron Spectroscopy and Related Phenomena* **82**, 23–29 (1996).

- [47] B. Da, H. Shinotsuka, H. Yoshikawa, Z. J. Ding, and S. Tanuma, “Extended Mermin method for calculating the electron inelastic mean free path”, *Physical Review Letters* **113**, 1–5 (2014).
- [48] I. Abril, R. Garcia-Molina, C. D. Denton, F. J. Pérez-Pérez, and N. R. Arista, “Dielectric description of wakes and stopping powers in solids”, *Physical Review A* **58**, 357–366 (1998).
- [49] S. Heredia-Avalos, R. Garcia-Molina, J. M. Fernández-Varea, and I. Abril, “Calculated energy loss of swift He, Li, B, and N ions in SiO₂, Al₂O₃, and ZrO₂”, *Physical Review A* **72**, 1–9 (2005).
- [50] I. Abril, C. D. Denton, P. de Vera, I. Kyriakou, D. Emfietzoglou, and R. Garcia-Molina, “Effect of the bethe surface description on the electronic excitations induced by energetic proton beams in liquid water and DNA”, *Nuclear Instruments and Methods in Physics Research, Section B: Beam Interactions with Materials and Atoms* **268**, 1763–1767 (2010).
- [51] U Fano, “Penetration of protons, alpha particles, and mesons”, *Annual Review of Nuclear Science* **13**, 1–66 (1963).
- [52] J. D. Bourke and C. T. Chantler, “Momentum-dependent lifetime broadening of electron energy loss spectra: a self-consistent coupled-plasmon model”, *Journal of Physical Chemistry Letters* **6**, 314–319 (2015).
- [53] D. C. Joy, “A database on electron-solid interactions”, *Scanning* **17**, 270–275 (1995).
- [54] F. Salvat-Pujol and J. S. Villarrubia, “Conventional vs. model-based measurement of patterned line widths from scanning electron microscopy profiles”, *Ultramicroscopy* **206**, 112819 (2019).
- [55] O. Y. Ridzel, V. Astašauskas, and W. S. Werner, “Low energy (1–100 eV) electron inelastic mean free path (IMFP) values determined from analysis of secondary electron yields (SEY) in the incident energy range of 0.1–10 keV”, *Journal of Electron Spectroscopy and Related Phenomena* **241**, 146824 (2020).
- [56] V. Astašauskas, A. Bellissimo, P. Kuksa, C. Tomastik, H. Kalbe, and W. S. Werner, “Optical and electronic properties of amorphous silicon dioxide by single and double electron spectroscopy”, *Journal of Electron Spectroscopy and Related Phenomena* **241**, 146829 (2020).
- [57] P. Zhang, “Monte Carlo study for correcting the broadened line-scan profile in scanning electron microscopy”, *Journal of Microscopy* **277**, 23–31 (2020).

- [58] J. P. Ganachaud and A. Mokrani, “Theoretical study of the secondary electron emission of insulating targets”, *Surface Science* **334**, 329–341 (1995).
- [59] S. T. Perkins, D. E. Cullen, and S. M. Seltzer, *Tables and graphs of electron-interaction cross-sections from 10 eV to 100 GeV derived from the LLNL evaluated electron data library (EEDL), Z=1–100*, tech. rep. UCRL-50400-Vol.31 (Lawrence Livermore National Laboratory, 1991).
- [60] D. A. Brown, M. B. Chadwick, R. Capote, A. C. Kahler, A. Trkov, M. W. Herman, A. A. Sonzogni, Y. Danon, A. D. Carlson, M. Dunn, *et al.*, “ENDF/B-VIII.0: the 8th major release of the nuclear reaction data library with CIELO-project cross sections, new standards and thermal scattering data”, *Nuclear Data Sheets* **148**, 1–142 (2018).
- [61] M. S. Chung and T. E. Everhart, “Role of plasmon decay in secondary electron emission in the nearly-free-electron metals. Application to aluminum”, *Physical Review B* **15**, 4699–4715 (1977).
- [62] A. Dubus, J.-C. Dehaes, J.-P. Ganachaud, A. Hafni, and M. Cailler, “Monte carlo evaluation of the influence of the interaction cross sections on the secondary-electron-emission yields from polycrystalline aluminum targets”, *Physical Review B* **47**, 11056 (1993).
- [63] Z. Ding, X. Tang, and R. Shimizu, “Monte carlo study of secondary electron emission”, *Journal of Applied Physics* **89**, 718–726 (2001).
- [64] S. F. Mao, Y. G. Li, R. G. Zeng, and Z. J. Ding, “Electron inelastic scattering and secondary electron emission calculated without the single pole approximation”, *Journal of Applied Physics* **104**, 114907 (2008).
- [65] S. Agostinelli, J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce, M. Asai, D. Axen, S. Banerjee, G. Barrand, *et al.*, “GEANT4—a simulation toolkit”, *Nuclear instruments and methods in physics research section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **506**, 250–303 (2003).
- [66] D. J. Griffiths, *Introduction to quantum mechanics* (Prentice Hall, 1995).
- [67] K. Murata, T. Matsukawa, and R. Shimizu, “Monte carlo calculations on electron scattering in a solid target”, *Japanese Journal of Applied Physics* **10**, 678–686 (1971).
- [68] A. Jablonski, “Elastic scattering and quantification in AES and XPS”, *Surface and Interface Analysis* **14**, 659–685 (1989).

- [69] N. F. Mott, “The scattering of fast electrons by atomic nuclei”, *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **124**, 425–442 (1929).
- [70] Z. Czyzewski, D. O. N. MacCallum, A. Romig, and D. C. Joy, “Calculations of Mott scattering cross section”, *Journal of Applied Physics* **68**, 3066–3072 (1990).
- [71] M. Dapor, “Elastic scattering calculations for electrons and positrons in solid targets”, *Journal of Applied Physics* **79**, 8406–8411 (1996).
- [72] F. Salvat, A. Jablonski, and C. J. Powell, “ELSEPA—Dirac partial-wave calculation of elastic scattering of electrons and positrons by atoms, positive ions and molecules”, *Computer Physics Communications* **165**, 157–190 (2005).
- [73] J. M. Ziman, *Electrons and phonons: the theory of transport phenomena in solids*, The International Series of Monographs on Physics (Oxford University Press, 1960).
- [74] A. Haug, *Theoretical solid state physics*, Vol. 2, International Series in Natural Philosophy (Pergamon Press, 1972).
- [75] J. Sólyom, *Fundamentals of the physics of solids: electronic properties*, Vol. 2 (Springer, 2009).
- [76] E. M. Conwell and M. O. Vassell, “High-field transport in n-type GaAs”, *Physical Review* **166**, 797–821 (1968).
- [77] J. N. Bradford and S. Woolf, “Electron-acoustic phonon scattering in SiO₂ determined from a pseudo-potential for energies of $E > E_{BZ}$ ”, *Journal of Applied Physics* **70**, 490–492 (1991).
- [78] E. Schreiber, “Monte-Carlo Simulationen des elektronischen Hochfeldtransportes in dielektrischen Schichten”, PhD thesis (Universität Rostock, 2001).
- [79] H.-J. Fitting, E. Schreiber, J.-C. Kuhr, and A. von Czarnowski, “Attenuation and escape depths of low-energy electron emission”, *Journal of Electron Spectroscopy and Related Phenomena* **119**, 35–47 (2001).
- [80] E. Schreiber and H.-J. Fitting, “Monte Carlo simulation of secondary electron emission from the insulator SiO₂”, *Journal of Electron Spectroscopy and Related Phenomena* **124**, 25–37 (2002).

- [81] E. Pop, R. W. Dutton, and K. E. Goodson, “Analytic band Monte Carlo model for electron transport in Si including acoustic and optical phonon dispersion”, *Journal of Applied Physics* **96**, 4998–5005 (2004).
- [82] A. Akkerman and M. Murat, “Electron-phonon interactions in silicon: mean free paths, related distributions and transport characteristics”, *Nuclear Instruments and Methods in Physics Research, Section B: Beam Interactions with Materials and Atoms* **350**, 49–54 (2015).
- [83] M. Sparks, D. L. Mills, R. Warren, T. Holstein, A. A. Maradudin, L. J. Sham, E. Loh, and D. F. King, “Theory of electron-avalanche breakdown in solids”, *Physical Review B* **24**, 3519–3536 (1981).
- [84] M. V. Fischetti and S. E. Laux, “Band structure, deformation potentials, and carrier mobility in strained Si, Ge, and SiGe alloys”, *Journal of Applied Physics* **80**, 2234–2252 (1996).
- [85] F. Mousty, P. Ostoja, and L. Passari, “Relationship between resistivity and phosphorus concentration in silicon”, *Journal of Applied Physics* **45**, 4576–4580 (1974).
- [86] R. N. Dexter, B. Lax, A. F. Kip, and G. Dresselhaus, “Effective masses of electrons in silicon”, *Physical Review* **96**, 222–223 (1954).
- [87] L. van Kessel and C. W. Hagen, “Nebula: Monte Carlo simulator of electron-matter interaction”, *SoftwareX* **12**, 100605 (2020).
- [88] P. Cizmar, A. E. Vladár, B. Ming, and M. T. Postek, “Simulated SEM images for resolution measurement”, *Scanning* **30**, 381–391 (2008).
- [89] S. Babin, S. S. Borisov, and V. P. Trifonenkov, “Fast analytical modeling of SEM images at a high level of accuracy”, *Proceedings of SPIE* **9424**, 94240I (2015).
- [90] T. Verduin, S. Lokhorst, and C. Hagen, “GPU accelerated Monte-Carlo simulation of SEM images for metrology”, *Proceedings of SPIE* **9778**, 97780D (2016).
- [91] T. Verduin, S. R. Lokhorst, P. Kruit, and C. W. Hagen, “The effect of sidewall roughness on line edge roughness in top-down scanning electron microscopy images”, *Proceedings of SPIE* **9424**, 942405 (2015).
- [92] T. Möller and B. Trumbore, “Fast, minimum storage ray-triangle intersection”, *Journal of Graphics Tools* **2**, 21–28 (1997).

- [93] A. L. dos Santos, V. Teichrieb, and J. Lindoso, “Review and comparative study of ray traversal algorithms on a modern GPU architecture”, 22nd International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision, 203–212 (2014).
- [94] Z.-J. Ding and R. Shimizu, “A Monte Carlo modeling of electron interaction with solids including cascade secondary electron production”, *Scanning* **18**, 92–113 (1996).
- [95] L. Reimer, *Scanning electron microscopy: physics of image formation and microanalysis*, Second edition (Springer, 1998).
- [96] K. T. Arat, J. Bolten, A. C. Zonneville, P. Kruit, and C. W. Hagen, “Estimating step heights from top-down SEM images”, *Microscopy and Microanalysis* **25**, 903–911 (2019).
- [97] I. Bojko, N. Hilleret, and C. Scheuerlein, “Influence of air exposures and thermal treatments on the secondary electron yield of copper”, *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films* **18**, 972–979 (2000).
- [98] N. Hilleret, C. Scheuerlein, and M. Tadorelli, “The secondary-electron yield of air-exposed metal surfaces”, *Applied Physics A: Materials Science and Processing* **76**, 1085–1091 (2003).
- [99] I. M. Bronstein and B. S. Fraiman, in *Vtorichnaya elektronnaya emissiya* (Nauka, Moskva, 1969), pp. 220, 340.
- [100] I. M. Bronshtein and R. B. Segal, “An investigation of electron reflection from certain metals”, *Soviet physics: solid state* **1**, 1142–1145 (1960).
- [101] I. M. Bronshtein and R. B. Segal, “Inelastic scattering of electrons and secondary-electron in certain metals I”, *Soviet physics: solid state* **1**, 1365–1496 (1960).
- [102] I. M. Bronshtein and B. S. Fraiman, “Secondary electron emission from some solid bodies”, *Soviet physics: solid state* **3**, 2087–2088 (1962).
- [103] I. M. Bronshtein and B. S. Fraiman, “Secondary emission properties of metals and semiconductors and the periodic system of elements”, *Soviet physics: solid state* **3**, 2337–2339 (1962).
- [104] I. M. Bronshtein and S. S. Denisov, “Effect of the work function on parameters of secondary electron emission”, *Soviet physics: solid state* **6**, 1515–1518 (1965).

- [105] I. M. Bronshtein and S. S. Denisov, “Inelastic scattering of electrons in solids with the primary beam at oblique incidence”, *Soviet physics: solid state* **6**, 2106–2110 (1965).
- [106] C. G. Walker, M. M. El-Gomati, A. M. Assa’d, and M. Zadražil, “The secondary electron emission yield for 24 solid elements excited by primary electrons in the range 250–5000 eV: a theory/experiment comparison”, *Scanning* **30**, 365–380 (2008).
- [107] M. M. El Gomati, C. G. Walker, A. M. Assa’d, and M. Zadražil, “Theory experiment comparison of the electron backscattering factor from solids at low electron energy (250–5,000 eV)”, *Scanning* **30**, 2–15 (2008).
- [108] L. van Kessel, C. W. Hagen, and P. Kruit, “Surface effects in simulations of scanning electron microscopy images”, *Journal of Micro/Nanolithography, MEMS, and MOEMS* **18**, 044002 (2019).
- [109] C. G. Frase, E. Buhr, and K. Dirscherl, “CD characterization of nanostructures in SEM metrology”, *Measurement Science and Technology* **18**, 510–519 (2007).
- [110] C. A. Mack and B. D. Bunday, “Improvements to the analytical linescan model for SEM metrology”, in *Proceedings of SPIE*, Vol. 9778, edited by M. I. Sanchez and V. A. Ukraintsev (2016), 97780A.
- [111] F. J. Pijper and P. Kruit, “Detection of energy-selected secondary electrons in coincidence with energy-loss events in thin carbon foils”, *Physical Review B* **44**, 9192–9200 (1991).
- [112] H. Mullejans, A. L. Bleloch, A. Howie, and C. A. Walsh, “Secondary-electron emission from magnesium oxide”, *Philosophical Magazine Letters* **68**, 145–151 (1993).
- [113] M. R. Scheinfein, J. Drucker, and J. K. Weiss, “Secondary-electron production pathways determined by coincidence electron spectroscopy”, *Physical Review B* **47**, 4068–4071 (1993).
- [114] W. S. M. Werner, A. Ruocco, F. Offi, S. Iacobucci, W. Smekal, H. Winter, and G. Stefani, “Role of surface and bulk plasmon decay in secondary electron emission”, *Physical Review B* **78**, 1–4 (2008).
- [115] F. J. García de Abajo, “Optical excitations in electron microscopy”, *Reviews of Modern Physics* **82**, 209–275 (2010).

- [116] F. Salvat-Pujol and W. S. M. Werner, "Surface excitations in electron spectroscopy. part I: dielectric formalism and Monte Carlo algorithm", *Surface and Interface Analysis* **45**, 873–894 (2013).
- [117] C. W. Peterson and B. W. Knight, "Causality calculations in the time domain: an efficient alternative to the Kramers-Kronig method", *Journal of the Optical Society of America* **63**, 1238–1242 (1973).
- [118] F. Salvat-Pujol, "Secondary-electron emission from solids: coincidence experiments and dielectric formalism", PhD thesis (TU Wien, 2012).
- [119] L. van Kessel, T. Huisman, and C. W. Hagen, "Understanding the influence of three-dimensional sidewall roughness on observed line-edge roughness in scanning electron microscopy images", *Journal of Micro/Nanolithography, MEMS, and MOEMS* **19**, 034002 (2020).
- [120] G. F. Lorusso, T. Sutani, V. Rutigliani, F. Van Roey, A. Moussa, A.-L. Charley, C. Mack, P. Naulleau, C. Perera, V. Constantoudis, M. Ikota, T. Ishimoto, and S. Koshihara, "Need for LWR metrology standardization: the imec roughness protocol", *Journal of Micro/Nanolithography, MEMS, and MOEMS* **17**, 041009 (2018).
- [121] C. A. Mack, "Systematic errors in the measurement of power spectral density", *Journal of Micro/Nanolithography, MEMS, and MOEMS* **12**, 033016 (2013).
- [122] A. Hiraiwa and A. Nishida, "Statistical- and image-noise effects on experimental spectrum of line-edge and line-width roughness", *Journal of Micro/Nanolithography, MEMS, and MOEMS* **9**, 041210 (2010).
- [123] D. L. Goldfarb, A. P. Mahorowala, G. M. Gallatin, K. E. Petrillo, K. Temple, M. Angelopoulos, S. Rasgon, H. H. Sawin, S. D. Allen, M. C. Lawson, and R. W. Kwong, "Effect of thin-film imaging on line edge roughness transfer to underlayers during etch processes", *Journal of Vacuum Science & Technology B* **22**, 647–653 (2004).
- [124] J. Foucher, A. L. Fabre, and P. Gautier, "CD-AFM versus CD-SEM for resist LER and LWR measurements", *Proceedings of SPIE* **6152**, 61520V (2006).
- [125] G. A. Dahlen, H.-C. Liu, M. Osborn, J. R. Osborne, B. Tracy, and A. Del Rosario, "TEM validation of CD AFM image reconstruction: part II", *Proceedings of SPIE* **6922**, 69220K (2008).
- [126] R. Kizu, I. Misumi, A. Hirai, and S. Gonda, "Accurate vertical sidewall measurement by a metrological tilting-AFM for reference metrology of line edge roughness", *Proceedings of SPIE* **10959**, 109592B (2019).

- [127] S. A. George, P. P. Naulleau, A. Krishnamoorthy, Z. Wu, E. W. Rutter Jr, J. T. Kennedy, S. Y. Xie, K. Y. Flanigan, and T. I. Wallow, "Characterization of line-edge roughness (LER) propagation from resists: underlayer interfaces in ultrathin resist films", *Proceedings of SPIE* **7636**, 763605 (2010).
- [128] V. Constantoudis, G. Kokkoris, E. Gogolides, E. Pargon, and M. Martin, "Effects of resist sidewall morphology on line-edge roughness reduction and transfer during etching: is the resist sidewall after development isotropic or anisotropic?", *Journal of Micro/Nanolithography, MEMS, and MOEMS* **9**, 041209 (2010).
- [129] J. Foucher, A. Pikon, C. Andes, and J. Thackeray, "Impact of acid diffusion length on resist LER and LWR measured by CD-AFM and CD-SEM", *Proceedings of SPIE* **6518**, 65181Q (2007).
- [130] M. Fouchier, E. Pargon, and B. Bardet, "An atomic force microscopy-based method for line edge roughness measurement", *Journal of Applied Physics* **113**, 104903 (2013).
- [131] M. Fouchier, E. Pargon, and B. Bardet, "A new method based on AFM for the study of photoresist sidewall smoothing and LER transfer during gate patterning", *Proceedings of SPIE* **8685**, 86850B (2013).
- [132] R. A. Lawson and C. L. Henderson, "Understanding the relationship between true and measured resist feature critical dimension and line edge roughness using a detailed scanning electron microscopy simulator", *Journal of Vacuum Science & Technology B* **28**, C6H34–C6H39 (2010).
- [133] R. A. Lawson and C. L. Henderson, "Investigating SEM metrology effects using a detailed SEM simulation and stochastic resist model", *Proceedings of SPIE* **9424**, 94240K (2015).
- [134] S. K. Sinha, E. B. Sirota, S. Garoff, and H. B. Stanley, "X-ray and neutron scattering from rough surfaces", *Physical Review B* **38**, 2297–2311 (1988).
- [135] G. Palasantzas, "Roughness spectrum and surface width of self-affine fractal surfaces via the K-correlation model", *Physical Review B* **48**, 14472 (1993).
- [136] E. I. Thorsos, "The validity of the Kirchhoff approximation for rough surface scattering using a Gaussian roughness spectrum", *The Journal of the Acoustical Society of America* **83**, 78–92 (1988).
- [137] C. A. Mack, "Generating random rough edges, surfaces, and volumes", *Applied Optics* **52**, 1472–1480 (2013).

- [138] T. Verduin, P. Kruit, and C. W. Hagen, “Determination of line edge roughness in low-dose top-down scanning electron microscopy images”, *Journal of Micro/Nanolithography, MEMS, and MOEMS* **13**, 033009 (2014).
- [139] C. A. Mack and B. D. Bunday, “Using the analytical linescan model for SEM metrology”, *Proceedings of SPIE* **10145**, 101451R (2017).
- [140] J. S. Villarrubia and B. D. Bunday, “Unbiased estimation of linewidth roughness”, *Proceedings of SPIE* **5752**, 480–488 (2005).
- [141] E. S. Putna, T. R. Younkin, M. Chandhok, and K. Frasure, “EUV lithography for 30nm half pitch and beyond: exploring resolution, sensitivity, and LWR tradeoffs”, *Proceedings of SPIE* **7273**, 72731L (2009).
- [142] L. van Look, J. Bekaert, A. Frommhold, E. Hendrickx, G. Rispens, and G. Schiffelers, “Optimization and stability of CD variability in pitch 40 nm contact holes on NXE:3300”, *Proceedings of SPIE* **10809**, 108090M (2018).
- [143] Y. B. Zou, M. S. S. Khan, H. M. Li, Y. G. Li, W. Li, S. T. Gao, L. S. Liu, and Z. J. Ding, “Use of model-based library in critical dimension measurement by CD-SEM”, *Measurement: Journal of the International Measurement Confederation* **123**, 150–162 (2018).
- [144] N. F. Mott and H. S. W. Massey, *The theory of atomic collisions*, third edition (Oxford University Press, Oxford, 1965).
- [145] C. J. Tung, T. L. Chou, and C. M. Kwei, “Stopping power of diamond for low energy electrons”, *Radiation Effects* **59**, 7–11 (1981).
- [146] V. I. Ochkur, “The Born-Oppenheimer method in the theory of atomic collisions”, *Soviet Physics, Journal of Experimental and Theoretical Physics* **18**, 503–508 (1964).
- [147] V. I. Ochkur, “Ionization of the hydrogen atom by electron impact with allowance for the exchange”, *Soviet Physics, Journal of Experimental and Theoretical Physics* **20**, 1175–1178 (1965).
- [148] J. M. Fernández-Varea, X. Llovet, and F. Salvat, “Cross sections for electron interactions in condensed matter”, *Surface and Interface Analysis* **37**, 824–832 (2005).

- [149] J. M. Fernández-Varea, F. Salvat, M. Dingfelder, and D. Liljequist, “A relativistic optical-data model for inelastic scattering of electrons and positrons in condensed matter”, *Nuclear Instruments and Methods in Physics Research, Section B: Beam Interactions with Materials and Atoms* **229**, 187–218 (2005).
- [150] J. M. Fernández-Varea, D. Liljequist, S. Csillag, R. Rätty, and F. Salvat, “Monte Carlo simulation of 0.1–100 keV electron and positron transport in solids using optical data and partial wave methods”, *Nuclear Instruments and Methods in Physics Research, Section B: Beam Interactions with Materials and Atoms* **108**, 35–50 (1996).
- [151] W. M. Saslow and G. F. Reiter, “Plasmons and characteristic energy loss in periodic solids”, *Physical Review B* **7**, 2995–3003 (1973).
- [152] C. Sevik and C. Bulutay, “Theoretical study of the insulating oxides and nitrides: SiO_2 , GeO_2 , Al_2O_3 , Si_3N_4 , and Ge_3N_4 ”, *Journal of Materials Science* **42**, 6555–6565 (2007).
- [153] J. Zhu, Z. An, M. Liu, and L. Tian, “Measurements of the K-shell ionization cross sections of Si by 3–25-keV electron impact using the thick-target method”, *Physical Review A* **79**, 1–7 (2009).
- [154] S. P. Limandri, M. A. Vasconcellos, R. Hinrichs, and J. C. Trincavelli, “Experimental determination of cross sections for K-shell ionization by electron impact for C, O, Al, Si, and Ti”, *Physical Review A* **86**, 1–10 (2012).
- [155] H. V. Rahangdale, M. Guerra, P. K. Das, S. De, J. P. Santos, D. Mitra, and S. Saha, “Determination of subshell-resolved *L*-shell-ionization cross sections of gold induced by 15–40-keV electrons”, *Physical Review A* **89**, 052708 (2014).
- [156] D. Davis, V. Mistry, and C. Quarles, “Inner shell ionization of copper, silver and gold by electron bombardment”, *Physics Letters A* **38**, 169–170 (1972).
- [157] A. Tonomura, J. Endo, T. Matsuda, T. Kawasaki, and H. Ezawa, “Demonstration of single-electron buildup of an interference pattern”, *American Journal of Physics* **57**, 117–120 (1989).
- [158] I. G. Austin and N. F. Mott, “Polarons in crystalline and non-crystalline materials”, *Advances in Physics* **18**, 41–102 (1969).

- [159] M. Dapor, “Secondary electron emission yield calculation performed using two different Monte Carlo strategies”, *Nuclear Instruments and Methods in Physics Research, Section B: Beam Interactions with Materials and Atoms* **269**, 1668–1671 (2011).
- [160] M. Dapor, “Monte Carlo simulation of secondary electron emission from dielectric targets”, *Journal of Physics: Conference Series* **402**, 012003 (2012).
- [161] K.J. Rietwyk, S.L. Wong, L. Cao, K.M. O’Donnell, L. Ley, A.T.S. Wee, and C.I. Pakes, “Work function and electron affinity of the fluorine-terminated (100) diamond surface”, *Applied Physics Letters* **102**, 091604 (2013).
- [162] K. Akada, S. Obata, and K. Saiki, “Work function lowering of graphite by sequential surface modifications: nitrogen and hydrogen plasma treatment”, *ACS Omega* **4**, 16531–16535 (2019).
- [163] J. Cazaux, “Mechanisms of charging in electron spectroscopy”, *Journal of Electron Spectroscopy and Related Phenomena* **105**, 155–185 (1999).
- [164] J. Cazaux, “About the mechanisms of charging in EPMA, SEM, and ESEM with their time evolution”, *Microscopy and Microanalysis* **10**, 670–684 (2004).
- [165] J. Cazaux, “Charging in scanning electron microscopy “from inside and outside””, *Scanning* **26**, 181–203 (2004).
- [166] Y.-U. Ko, M.-S. Chung, and S.-W. Kim, “Monte Carlo simulation of charging effects on linewidth metrology”, *Scanning* **20**, 447–455 (1998).
- [167] L. Grella, G. Lorusso, and D.L. Adler, “Simulations of scanning electron microscopy imaging and charging of insulating structures”, *Scanning* **25**, 300–308 (2003).
- [168] K. Ohya, K. Inai, H. Kuwada, T. Hayashi, and M. Saito, “Dynamic simulation of secondary electron emission and charging up of an insulating material”, *Surface and Coatings Technology* **202**, 5310–5313 (2008).
- [169] C. Li, S.F. Mao, Y.B. Zou, Y.G. Li, P. Zhang, H.M. Li, and Z.J. Ding, “A Monte Carlo modeling on charging effect for structures with arbitrary geometries”, *Journal of Physics D: Applied Physics* **51**, 165301 (2018).

- [170] B. Raftari, N. Budko, and K. Vuik, “A modified and calibrated drift-diffusion-reaction model for time-domain analysis of charging phenomena in electron-beam irradiated insulators”, *AIP Advances* **8**, 015307 (2018).
- [171] B. E. Weir, C. C. Leung, P. J. Silverman, and M. A. Alam, “Gate dielectric breakdown in the time-scale of ESD events”, *Microelectronics Reliability* **45**, 427–436 (2005).
- [172] R. F. Egerton, P. Li, and M. Malac, “Radiation damage in the TEM and SEM”, *Micron* **35**, 399–409 (2004).
- [173] V. Ageev, “Desorption induced by electronic transitions”, *Progress in Surface Science* **47**, 55–203 (1994).
- [174] T. V. Pistor, T. I. Wallow, C. N. Anderson, and P. P. Naulleau, “Photoresist shrinkage effects at EUV”, *Proceedings of SPIE* **7969**, 796917 (2011).
- [175] B. Bunday, A. Cordes, C. Hartig, J. Allgair, A. Vaid, E. Solecky, and N. Rana, “Time-dependent electron-beam-induced photoresist shrinkage effects”, *Journal of Micro/Nanolithography, MEMS, and MOEMS* **11**, 023007–1 (2012).
- [176] R. Hamming, *Numerical methods for scientists and engineers*, 2nd edition (Dover Publications, 1973).
- [177] J. Rumble, *CRC handbook of chemistry and physics*, 98th Edition (Taylor & Francis Group, 2017).
- [178] K. Ohya, A. Chen, T. Tanabe, J. Kawata, K. Nishimura, D. Kato, and T. Kato, *ELECTRAN–Monte Carlo program of secondary electron emission from monoatomic solids under the impact of 0.1–10 keV electrons*, tech. rep. NIFS-DATA-84 (National Institute for Fusion Science, 2004).
- [179] S. S. Nekrashevich and V. A. Gritsenko, “Electronic structure of silicon dioxide (a review)”, *Physics of the Solid State* **56**, 207–222 (2014).
- [180] C. Kittel, *Introduction to solid state physics*, Eighth edition (John Wiley & Sons, 2005).
- [181] J. A. Hagen, W. Li, A. J. Steckl, and J. G. Grote, “Enhanced emission efficiency in organic light-emitting diodes using deoxyribonucleic acid complex as an electron blocking layer”, *Applied Physics Letters* **88**, 171109 (2006).
- [182] D. V. Melnikov and J. R. Chelikowsky, “Electron affinities and ionization energies in si and ge nanocrystals”, *Physical Review B* **69**, 113305 (2004).

- [183] N. Fujimura, A. Ohta, K. Makihara, and S. Miyazaki, "Evaluation of valence band top and electron affinity of SiO_2 and Si-based semiconductors using X-ray photoelectron spectroscopy", *Japanese Journal of Applied Physics* **55**, 08PC06 (2016).
- [184] J. W. Jewett and R. A. Serway, *Physics for scientists and engineers with modern physics*, 8th international edition (Cengage Learning, 2010).
- [185] S. Suzuki, C. Bower, T. Kiyokura, K. G. Nath, Y. Watanabe, and O. Zhou, "Photoemission spectroscopy of single-walled carbon nanotube bundles", *Journal of Electron Spectroscopy and Related Phenomena* **114**, 225–228 (2001).
- [186] P. Becker, P. Scyfried, and H. Siegert, "The lattice parameter of highly pure silicon single crystals", *Zeitschrift für Physik B Condensed Matter* **48**, 17–21 (1982).
- [187] W. Witt, "Absolute präzisionsbestimmung von gitterkonstanten an germanium-und aluminium-einkristallen mit elektroneninterferenzen", *Zeitschrift für Naturforschung A* **22**, 92–95 (1967).
- [188] M. E. Straumanis and L. S. Yu, "Lattice parameters, densities, expansion coefficients and perfection of structure of Cu and of Cu-In α phase", *Acta Crystallographica Section A: Crystal Physics, Diffraction, Theoretical and General Crystallography* **25**, 676–682 (1969).
- [189] P. Trucano and R. Chen, "Structure of graphite by neutron diffraction", *Nature* **258**, 136–137 (1975).
- [190] O. Madelung, U. Rössler, and M. Schulz, eds., *Silicon (Si) sound velocities: datasheet from Landolt-Börnstein - group III condensed matter volume 41a1a*.
- [191] D. R. Christman, *Dynamic properties of poly(methylmethacrylate) (PMMA) (plexiglas)*, tech. rep. AD743547 (General Motors Corporation, 1972).
- [192] H. J. McSkimin, "Measurement of elastic constants at low temperatures by means of ultrasonic waves—data for silicon and germanium single crystals, and for fused silica", *Journal of applied physics* **24**, 988–997 (1953).
- [193] D. M. Zhernokletov, V. V. Milyavskiy, K. V. Khishchenko, A. A. Charakhchyan, T. I. Borodina, G. E. Val'vano, and A. Z. Zhuk, "Shock-wave loading of graphite in steel targets with conic cavities", in *AIP conference proceedings*, Vol. 955, 1 (2007), pp. 216–219.

-
- [194] O. Madelung, U. Rössler, and M. Schulz, eds., *Silicon (Si), deformation potentials: datasheet from Landolt-Börnstein - group III condensed matter volume 41a1β*.
- [195] K. Sugihara, "Thermoelectric power of graphite intercalation compounds", *Physical Review B* **28**, 2157 (1983).
- [196] M. Schowalter, J. T. Titantah, D. Lamoen, and P. Kruse, "Ab initio computation of the mean inner Coulomb potential of amorphous carbon structures", *Applied Physics Letters* **86**, 112102 (2005).