

Vegard Hennestad

Electromagnetic Cascades in Active Galactic Nuclei Cores

Master's thesis in Applied Physics and Mathematics

Supervisor: Michael Kachelrieß

June 2022

NTNU
Norwegian University of Science and Technology
Faculty of Natural Sciences
Department of Physics



Norwegian University of
Science and Technology

Vegard Hennestad

Electromagnetic Cascades in Active Galactic Nuclei Cores

Master's thesis in Applied Physics and Mathematics
Supervisor: Michael Kachelrieß
June 2022

Norwegian University of Science and Technology
Faculty of Natural Sciences
Department of Physics

Abstract

The following presents a program for the simulation of electromagnetic cascades, consisting of photons, electrons, and positrons, on the thermal background radiation from an accretion disk in an active galactic nucleus core. The program features an adaptive propagation algorithm which iteratively moves each particle a certain step length until an interaction with the background occurs, determined by comparing the total interaction probability at each step with a random number. Once an interaction occurs, the interaction parameters are determined by random sampling from their respective probability distributions. Electrons produced below a certain threshold energy are discarded to avoid excessive production of soft photons. The simulation agrees with optical depth calculations performed on the same system.

Preface

I have been interested in the natural sciences for as long as I can remember. But I can also clearly remember the two moments that violently reshaped my entire worldview and gave me a nearly spiritual perspective on the laws and awes of nature.

The first was upon reading Richard Dawkins' 1976 book *The Selfish Gene* in the eleventh grade. And although I would love to delve deep into the awesome explicative power of natural selection, a topic with the potential to tear so deeply into the philosophical questions about why you and I even exist that Dawkins has spent most of his life arguing with people who simply refuse to believe in it, it appears that my master's thesis is on the topic of physics, and not on evolutionary biology. But you need not worry that I have picked the wrong field of study, because physics has the potential to go even deeper.

Although the planets and stars would keep hurling about even if there was no life in the universe, it would not be as interesting to have physical laws if no one were around to experience them. Right now, we may be living at the first moment in the history of the entire universe that the Standard Model of particle physics has ever been hypothesized. And the mathematical complexity of quantum field theory would never have been able to confuse me if no one had discovered it first. So, it appears that there is some connection between life and physics, and this leads to the second moment to fully ignite my passion for science.

For while life is necessary to appreciate physics, it was on my voyage through the 1980 television series *Cosmos* that Carl Sagan showed me so clearly why physics is necessary to appreciate life. Biology can merely answer how life works and why it exists, but if one should wish to completely understand their brief presence in this world, they inevitably need to ask how the universe works and why it exists, too. Science is currently dealing with the how, and it has been doing so quite efficiently in the last couple of hundred years. We now have a pretty plausible explanation for how the universe put together all the necessary ingredients for life on our very planet. Then evolutionary biology explains how life occurred, and how a small strand of that life evolved into humans. And those humans turned out to be pretty curious about stuff.

What follows in this thesis is my contribution to the collective human curiosity. While it does not answer any of life's greatest questions, it does leave me feeling content with the five years I have spent achieving the scientific liter-

acy needed to write something like this. It has given me tremendous pleasure and a sense of belonging to finally feel that I understand the general picture of everything we know about the universe so far. While *Cosmos* can take a lot of credit for bringing me here, by illuminating all the major discoveries that led us to the scientific understanding we have today, it also lends some support to the idea that human consciousness and physics are intricately connected. For in the words of Carl Sagan: "We are a way for the Cosmos to know itself."

Acknowledgements

While I would not say that writing a master's thesis is harder than I imagined (because I have simply imagined it to be impossible), I can certainly say that it was the hardest thing I have ever done. I thank Michael Kachelrieß for his excellent guidance and everlasting patience.

I also thank my friends, who never faltered in their attempts to distract me from finishing my degree. I am looking forward to all the years to come, when such efforts are no longer needed.

Contents

Abstract	v
Preface	vii
Acknowledgements	ix
1 Introduction	1
1.1 Cascades in Extragalactic Background Light	1
1.2 Cascades in Active Galactic Nuclei	2
2 Theory	5
2.1 Thermal Radiation	5
2.1.1 Spectral Intensity	5
2.1.2 Spectral Photon Density	7
2.1.3 The Planck Distribution	8
2.2 Interaction Rate	9
2.2.1 Interaction Rate in Anisotropic Background	9
2.2.2 Pair Production	11
2.2.3 Inverse Compton Scattering	13
2.2.4 Interaction Probability	15
2.3 The Accretion Disk	17
3 The Program	21
3.1 Integration	21
3.1.1 The Integrand	22
3.1.2 Locating the Inner Integrand	22
3.1.3 Locating the Peak of the Inner Integrand	24
3.1.4 Locating the Left Limit of the Inner Integrand	25
3.1.5 Estimating the Inner Integral	27
3.1.6 Estimating the Outer Integral	28
3.1.7 Sampling the Angle and Frequency	29
3.2 Propagation	30
3.2.1 Step Size	30
3.2.2 Interactions	31

3.2.3 Storing Particles	32
3.3 Program structure	34
4 Results and Discussion	37
5 Conclusion	43
Bibliography	45

Chapter 1

Introduction

When looking up at the night sky, we see photons that have travelled trillions of miles to reach our eyes. The light we observe in the visible spectrum is only a small fraction of the vast number of particles that are launched into space from distant cosmic sources. On their years-long journey to reach the lens of your eye, they travel through an ocean of cosmic background radiation. Given the right circumstances, interactions can occur, so many of the particles that reach us are actually created during a journey through empty space.

Cascades occur when the resulting particles from an interaction can interact further, leading to a myriad of interactions which exponentially increases the amount of particles present. This thesis will consider such a cascade occurring on the background radiation from the accretion disk of an active galactic nucleus (AGN). It is inspired by the program ELMAG [1], which simulates electromagnetic cascades on the extragalactic background light (EBL). The processes involved in these electromagnetic cascades are illustrated in Fig. 1.1.

1.1 Cascades in Extragalactic Background Light

When an ultra-high-energy photon travels through the EBL, it will interact with background photons to produce an electron-positron pair. Hence, we cannot observe ultra-high-energy photons from distant sources directly, since they will undergo pair production long before they reach us. But we can observe them indirectly, by looking at the resulting particles from their interactions.

The electrons and positrons (from now on referred to collectively as electrons) can also interact with the background, through the process of inverse Compton scattering. While Compton scattering concerns high-energy photons hitting stationary electrons, inverse Compton scattering considers high-energy electrons interacting with low-energy photons. The electron transfers some of its energy to the photon, resulting in another high-energy photon that can further contribute to the cascade.

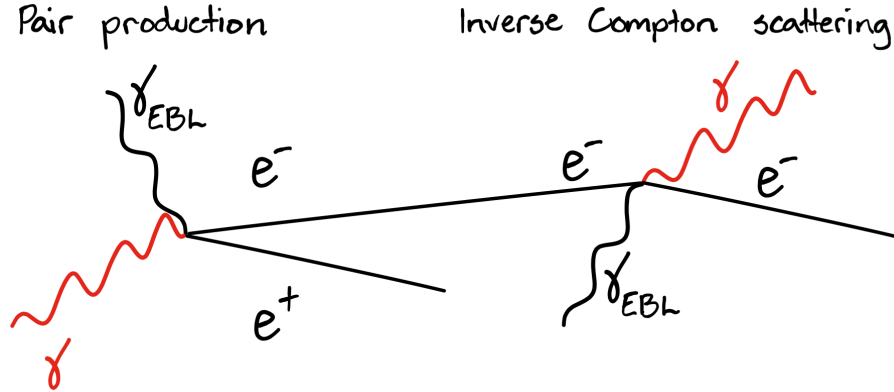


Figure 1.1: To the left we see an incoming high-energy photon, in red, interact with a low-energy background photon, creating an electron and a positron. To the right we see the electron interact with a background photon. The electron transfers some of its energy to the photon, resulting in yet another high-energy photon.

1.2 Cascades in Active Galactic Nuclei

Active galactic nuclei may be one of the most epic phenomena in our universe. At the center is a supermassive black hole, which is an unfathomable physical construct in and of itself. Matter falling into it forms an accretion disk, which emits immense quantities of electromagnetic energy as thermal radiation. The power output of quasars, the most powerful of AGNs, can be larger than that of entire galaxies [2]. AGNs perplexed scientists upon discovery, since their brightness would necessitate an incredible power output if they were truly as far away from Earth as their redshifts indicated. The current understanding of AGNs, that they are fueled by accretion onto a supermassive black hole, first appeared in the middle of the 1960s, just a few years after quasars were first discovered [3]. The source of the energy is the loss of potential energy in the gravitational field as matter falls inwards. Some AGNs have a pair of jets of high energy particles which extend in opposite directions. An AGN core is illustrated in Fig. 1.2.

Simulating an electromagnetic cascade near an AGN core presents a few challenges compared with cascades on the EBL. Firstly, while the EBL can be assumed to be isotropic, we cannot make such an assumption when the source of the background radiation is an accretion disk. We will, however, assume cylindrical symmetry, which means that any high-energy particle we are tracking needs to travel along the axis of symmetry of the system. This is not a terrible assumption, since jets of high-energy particles in AGNs are produced roughly in that direction. The difference in interaction rates and interaction parameters are assumed to be small for tiny angular deflections from this axis.

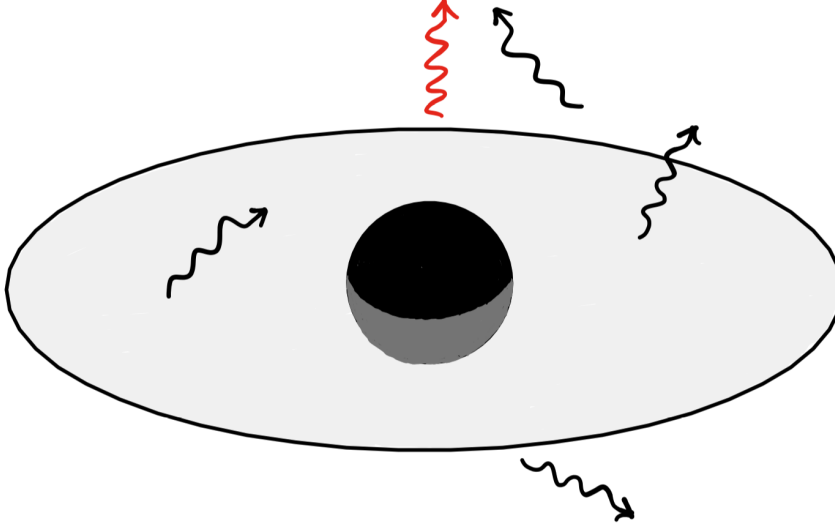


Figure 1.2: A simplified model of an AGN core. A supermassive black hole is surrounded by an accretion disk. The accretion disk radiates thermal photons, shown in black. We will consider high-energy particles moving away from the AGN core along the axis of symmetry.

The anisotropy does, however, necessitate a way of calculating the interaction rate for anisotropic backgrounds. This was done in the project work on the related topic of optical depth in AGN cores, which this thesis is a continuation of. Since it is considered essential to the background theory of the master's thesis, this derivation will be presented in its entirety in Ch. 2.

The second challenge is that the background radiation stemming from the accretion disk is not uniform, an assumption we can also make for the EBL. As particles are assumed to be moving along the axis of symmetry, we only need to introduce one spatial coordinate, namely the height above the center of the system. Since the interaction rate now changes as the particle moves, the probability that the particle interacts at a certain height no longer follows an exact exponential distribution, like it would in the uniform case. To overcome this, we will iteratively propagate each particle over small enough distances such that the probability of interaction in each step is small.

Chapter 2

Theory

Simulating an electromagnetic cascade requires mathematical expressions for the thermal background radiation and for the interaction rates of high-energy photons and electrons moving through this background.

2.1 Thermal Radiation

In order to calculate the rate of interaction with background radiation, the spectral photon density of the background radiation is needed. This is the number of photons per unit volume per unit frequency, and is given by

$$n_\nu = \frac{dn}{d\nu} = \frac{dN}{dV d\nu}, \quad (2.1)$$

where N is the number of background photons, V is a spatial volume, and ν is the frequency of the photons. The most important source of background radiation in an active galactic nucleus is the accretion disk surrounding the supermassive black hole. The electromagnetic radiation from the accretion disk will be treated as black-body. Since the background radiation surrounding the accretion disk is not isotropic, it is necessary to find an expression for the photon density which does not possess this restriction. Hence, a general expression for the spectral photon density away from a source will be derived in the following section.

2.1.1 Spectral Intensity

A thermal source of radiation has the spectral intensity I_ν , which is the radiated energy per unit time per unit projected area per unit solid angle per unit frequency. The projected area is illustrated in Fig. 2.1. The spectral intensity is given mathematically by

$$I_\nu = \frac{dE}{dt dA_\perp d\Omega_S d\nu}, \quad (2.2)$$

where E is the radiated energy, A_{\perp} is the projection of an area A on the source onto a plane perpendicular to the direction in which it radiates, Ω_S is the solid angle element into which it radiates, and t is time.




Figure 2.1: The projected area dA_{\perp} is the apparent area of dA when seen from an observer located in the solid angle element $d\Omega_S$.

For a thermal source at rest, the spectral intensity is independent of direction. This means that the power per unit projected area in a given frequency interval which is radiated into a solid angle element is constant, independent of where the solid angle element is located.

We will now consider an observer which receives radiation from the source on an area element dS , shown in Fig. 2.2. We can define the analogous received spectral intensity as

$$I_{\nu}^R = \frac{dE}{dt dS_{\perp} d\Omega_R d\nu}, \quad (2.3)$$

where S_{\perp} is the projection of S and Ω_R is the solid angle of A when seen from the observer.

From conservation of energy, we know that the spectral power received at dS from $d\Omega_R$ must be the same as that emitted from dA into $d\Omega_S$,

$$I_{\nu} dA_{\perp} d\Omega_S = \frac{dE}{dt d\nu} = I_{\nu}^R dS_{\perp} d\Omega_R. \quad (2.4)$$

However, rearranging yields that

$$I_{\nu} \frac{dA_{\perp}}{d\Omega_S} = I_{\nu}^R \frac{dS_{\perp}}{d\Omega_R}. \quad (2.5)$$




Figure 2.2: The receiver area dS , with orthogonal projection dS_{\perp} , and the solid angle projection $d\Omega_R$ of the source area dA .

Using the definition of the steradian, we have that $dA_{\perp}/d\Omega_R = d^2 = dS_{\perp}/d\Omega_S$, where d is the distance between the source and the observer. The result is that the spectral intensity received by the observer is identical to the spectral intensity radiated by the source, and it is also independent of the distance between source and observer. Hence, the expected $1/d^2$ attenuation in received spectral power is caused by the reduction of the solid angle projection of the source as seen from the observer, not by a reduction of the intensity.

2.1.2 Spectral Photon Density

The spectral photon density is tied to the spectral energy density, the energy per unit volume per unit frequency, through the quantization of photon energy, $E = h\nu$, where h is the Planck constant. In order to express the former in the vicinity of the observer, we will derive the latter using the spectral intensity of the source. As we saw in the previous subsection, the thermal spectral intensity is independent of distance and direction, and this makes it a quantity we can use to find the desired energy density.

Since radiation moves at the speed of light c , the spectral energy density near the observer can be found by considering the energy from an angle element $d\Omega_R$ which hits the surface dS in time dt . This energy must be contained in the volume $c dt dS_{\perp}$, as shown in Fig. 2.3. Hence, we get the spectral energy

density at the observer by integrating over all solid angles $d\Omega_R$, yielding

$$u_\nu = \frac{dE}{dV d\nu} = \frac{1}{c} \int d\Omega_R \frac{dE}{dt dS_\perp d\Omega_R d\nu} = \frac{1}{c} \int d\Omega_R I_\nu, \quad (2.6)$$

where we have used Eq. (2.3) to tie the quantity to the spectral intensity. If the source is uniform and extends in all directions surrounding the observer, this leads to the familiar energy density of $4\pi I_\nu/c$.

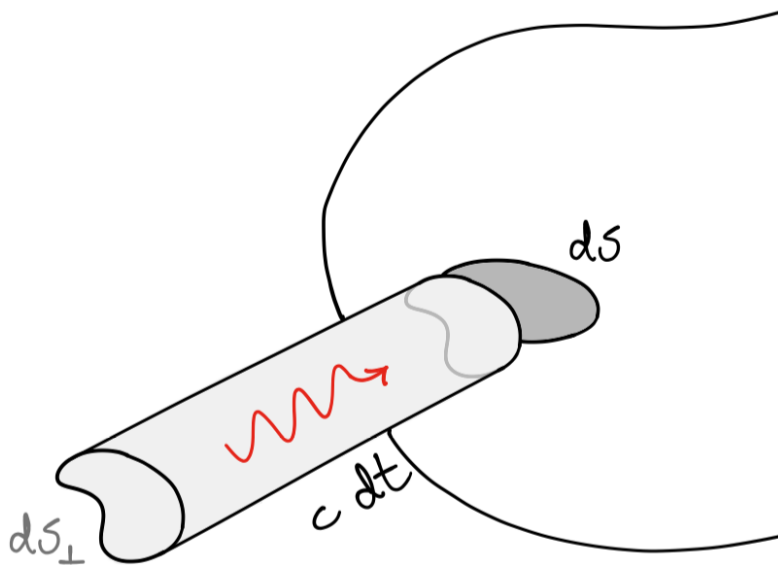


Figure 2.3: The volume containing the energy from $d\Omega_R$ which will hit the projection of the observer, and therefore the observer, in time interval dt .

By considering that the energy of a given photon is equal to $h\nu$, we obtain the relation $u_\nu = h\nu n_\nu$. Hence, the general expression for the spectral photon density becomes

$$n_\nu = \frac{1}{ch\nu} \int d\Omega_R I_\nu. \quad (2.7)$$

2.1.3 The Planck Distribution

A black body radiates with the spectral intensity B_ν , which is given by the Planck distribution:

$$B_\nu = \frac{2h}{c^2} \frac{\nu^3}{\exp(h\nu/kT) - 1}, \quad (2.8)$$

where T is the temperature of the object and k is the Boltzmann constant [4, p. 132]. The spectral intensity of an object at $T = 100$ K is shown in Fig. 2.4. It has a maximum of 1.90×10^{-10} erg s $^{-1}$ cm $^{-2}$ sr $^{-1}$ Hz $^{-1}$ at 5.9×10^{12} Hz.

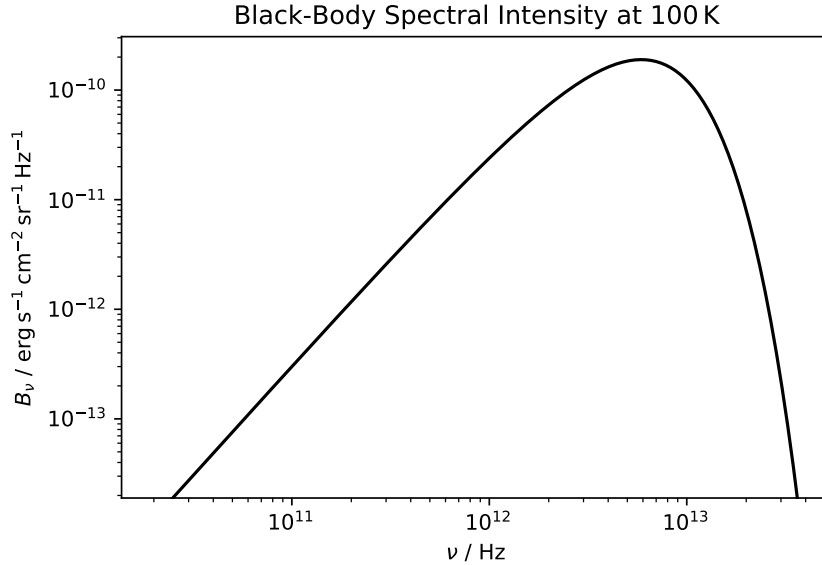


Figure 2.4: The Planck distribution for an object radiating at a temperature of 100 K.

2.2 Interaction Rate

We now consider a particle moving through the thermal background radiation. To model a cascade, we need a measure for how likely it is that an interaction happens as the particle travels a certain distance. This is quantified in the interaction rate, which in a uniform background is the inverse of the mean free path of the particle. Since the background in our model is not uniform, we need to evaluate the interaction rate for different values of the distance z from the center of the black hole.

The general expression for the interaction rate in an isotropic background of photons is given by

$$\mathcal{R} = \frac{1}{2\beta} \int d\nu n_\nu \int d\mu (1 - \beta\mu) \sigma, \quad (2.9)$$

where βc is the speed of the particle, $\mu = \cos\theta$, with θ the angle between the incident particles, and σ is the cross section of the interaction [5, p. 479]. We will use the $(1 - 1 - 1 - 1)$ metric. Four-momentum is denoted by p , while \vec{p} is three-momentum.

2.2.1 Interaction Rate in Anisotropic Background

Since Eq. (2.9) assumes an isotropic photon density, we will now rewrite it in such a way that it is compatible with the directional photon density derived in

Sec. 2.1.2.

By pulling out a factor of $1/2\pi$, we can reinstate the azimuthal angle into the inner integral. This makes it a solid angle integral over Ω , which represents the relative direction of the two incoming particles. Additionally, we replace the spectral photon density with a directional variant on the form of Eq. (2.7), yielding

$$\mathcal{R} = \frac{1}{4\pi\beta} \int d\nu \int d\Omega_{\text{R}} \frac{dn_{\nu}}{d\Omega_{\text{R}}} \int d\Omega (1 - \beta\mu) \sigma. \quad (2.10)$$

By letting $\theta = 0$ in the direction of the high-energy particle, Fig. 2.5 illustrates that the solid angle elements Ω and Ω_{R} can be directly tied together, as they are pointing in opposite directions. Let the angles Θ and Φ represent the polar and azimuthal angles in the coordinate system of Ω_{R} , which points towards the source. If $\Theta = 0$ aligns with $\theta = 0$, the angles are related through the transformation $\Theta \rightarrow \theta = \pi - \Theta$ and $\Phi \rightarrow \phi = \Phi + \pi$. The solid angle element $d\Omega_{\text{R}} = \sin\Theta d\Phi d\Theta$ is unchanged under this transformation, since $\sin\Theta = \sin(\pi - \Theta)$. Hence, we may safely rewrite the directional photon density as

$$\frac{dn_{\nu}}{d\Omega} = \frac{1}{ch\nu} I_{\nu}, \quad (2.11)$$

where the solid angle element now is defined to be pointing away from the source, not towards it.




Figure 2.5: The angle between the incident photons and its relationship with the solid angle projection $d\Omega_{\text{R}}$ of the source of the background radiation.

Redefining the directional spectral photon density lets us integrate away the solid angle Ω_R , which cancels the factor $1/4\pi$. Hence, the final expression for the interaction rate in an anisotropic background becomes

$$\mathcal{R} = \frac{1}{\beta} \int d\Omega \int d\nu \frac{dn_\nu}{d\Omega} (1 - \beta\mu) \sigma. \quad (2.12)$$

By inserting the Planck distribution into the directional spectral photon density, we obtain

$$\frac{dn_\nu}{d\Omega} = \frac{2}{c^3} \frac{\nu^2}{\exp(h\nu/kT) - 1} \quad (2.13)$$

for black-body radiation.

2.2.2 Pair Production

A pair of photons can produce charged particle-antiparticle pairs. Using the theory of quantum electrodynamics (QED), the relativistic quantum field theory which governs all electromagnetic interactions, one may compute the cross section for pair production, as well as Compton scattering. Since the electron is the lightest charged particle, the most common interaction is the creation of an electron-positron pair (from now on this is what we mean by pair production). The Feynman diagram for this interaction is shown in Fig. 2.6.




Figure 2.6: Two photons interact to create an electron-positron pair. This figure shows the t-channel, the u-channel is omitted.

The fundamental constraint for pair production is that there needs to be enough energy present to create the rest mass of the resulting particles. Since the energy of a particle depends on the frame of reference, we seek to express this in a way that is Lorentz invariant. We identify the center of momentum frame as the crucial frame of reference, because producing particles at rest in this frame must be the interaction requiring the least amount of energy in any frame. This leads us to the first Mandelstam variable, s . It is the squared center of momentum energy, so, for a pair of electrons produced at rest, we know that $\sqrt{s} = 2m_e c^2$, where m_e is the electron mass. The energies at which pair production can happen must therefore satisfy

$$s > 4m_e^2 c^4. \quad (2.14)$$

The pair production cross section is given by

$$\sigma_p = \frac{3}{4} \sigma_{\text{Th}} \frac{m_e^2 c^4}{s} \left[(3 - \beta_e^4) \ln \frac{1 + \beta_e}{1 - \beta_e} - 2\beta_e (2 - \beta_e^2) \right], \quad (2.15)$$

where

$$\sigma_{\text{Th}} = \frac{8\pi\alpha^2 \hbar^2}{3m_e^2 c^2} \quad (2.16)$$

is the Thomson cross section, with α the fine structure constant and \hbar the reduced Planck constant, and $\beta_e c$ is the speed of the resulting electrons in the center of momentum frame [1, p. 4].

We can find the squared center of momentum energy from the four-momenta of the incoming particles,

$$s = (p_1 + p_2)^2 c^2 = \left[\begin{pmatrix} E/c \\ E/c \\ 0 \\ 0 \end{pmatrix} + \begin{pmatrix} h\nu/c \\ h\nu \cos \theta/c \\ 0 \\ h\nu \sin \theta/c \end{pmatrix} \right]^2 c^2 = 2Eh\nu(1 - \mu). \quad (2.17)$$

Four-momenta are labelled according to Fig. 2.6. From conservation of energy and momentum, we know that $p_1 + p_2 = p_3 + p_4$. Hence,

$$s = (p_3 + p_4)^2 c^2 = \left[\begin{pmatrix} \sqrt{m_e^2 c^2 + \vec{p}_3^2} \\ \vec{p}_3 \end{pmatrix} + \begin{pmatrix} \sqrt{m_e^2 c^2 + \vec{p}_4^2} \\ \vec{p}_4 \end{pmatrix} \right]^2 c^2. \quad (2.18)$$

By considering the center of momentum frame, the sum $\vec{p}_3 + \vec{p}_4$ vanishes, since they are equal and opposite. We may rewrite the energy components of the four-momenta, which are identical in this frame, through the definition of relativistic energy,

$$E_3 = E_4 = \gamma m_e c^2, \quad (2.19)$$

where γ is the Lorentz factor $1/\sqrt{1 - \beta_e^2}$. Doing so, we obtain

$$s = (2\gamma m_e c^2)^2 = \frac{4m_e^2 c^4}{1 - \beta_e^2}. \quad (2.20)$$

Finally, we may solve for β_e , yielding

$$\beta_e = \sqrt{1 - \frac{4m_e^2 c^4}{s}}. \quad (2.21)$$

By inserting Eqs. (2.13), (2.15), (2.16), and (2.17) into Eq. (2.12), and using that $\beta = 1$ for the incoming photon, we arrive at the interaction rate

$$\mathcal{R}_p = \frac{\alpha^2 h}{2\pi c} \frac{1}{E} \int d\Omega \int d\nu \frac{\nu}{\exp(h\nu/kT) - 1} f_p(\beta_e), \quad (2.22)$$

where we have introduced the auxiliary function

$$f_p(\beta_e) = (3 - \beta_e^4) \ln \frac{1 + \beta_e}{1 - \beta_e} - 2\beta_e (2 - \beta_e^2). \quad (2.23)$$

2.2.3 Inverse Compton Scattering

The Feynman diagram for Compton scattering is shown in Fig. 2.7. Unlike in pair production, there is no minimal squared center of momentum energy required for an interaction to occur. This makes sense, since there are no new massive particles produced. In the low-energy regime, the interaction is known as Thomson scattering.




Figure 2.7: Scattering of electrons and photons. This figure shows the t-channel, the s-channel is omitted.

The Compton scattering cross section is given by

$$\sigma_C = \sigma_{\text{Th}} y \cdot \frac{3}{4} \left[\frac{-\ln y}{1-y} \left(1 - \frac{4y(1+y)}{(1-y)^2} \right) + \frac{8y}{(1-y)^2} + \frac{1+y}{2} \right], \quad (2.24)$$

where

$$y = \frac{m_e^2 c^4}{s} \quad (2.25)$$

is the minimal energy fraction taken by the outgoing electron [1, p. 4].

Once again, we find the squared center of momentum energy,

$$s = (p_1 + p_2)^2 c^2 = \left[\begin{pmatrix} \sqrt{m_e^2 c^2 + |\vec{p}_1|^2} \\ |\vec{p}_1| \\ 0 \\ 0 \end{pmatrix} + \begin{pmatrix} h\nu/c \\ h\nu \cos \theta/c \\ 0 \\ h\nu \sin \theta/c \end{pmatrix} \right]^2 c^2. \quad (2.26)$$

Four-momenta are labelled according to Fig. 2.7. By the definition of relativistic momentum, we may write

$$|\vec{p}_1| = \gamma m_e |\vec{v}_1| = \gamma m_e c^2 \frac{\beta}{c} = E \frac{\beta}{c}, \quad (2.27)$$

where \vec{v}_1 is the velocity of the incoming electron. Hence, evaluating Eq. (2.26), we get

$$s = m_e^2 c^4 + 2Eh\nu(1 - \beta\mu). \quad (2.28)$$

From the expression

$$E = \gamma m_e c^2 = \frac{m_e c^2}{\sqrt{1 - \beta^2}}, \quad (2.29)$$

we find that in terms of E , β is given by

$$\beta = \sqrt{1 - \frac{m_e^2 c^4}{E^2}}. \quad (2.30)$$

By inserting Eqs. (2.13), (2.16), and (2.24) into Eq. (2.12), we arrive at the interaction rate

$$\mathcal{R}_C = \frac{4\alpha^2 h^2}{3\pi m_e^2 c^5} \frac{1}{\beta} \int d\Omega \int d\nu \frac{\nu^2}{\exp(h\nu/kT) - 1} (1 - \beta\mu) f_C(y), \quad (2.31)$$

where we have introduced the auxiliary function

$$f_C(y) = y \cdot \frac{3}{4} \left[\frac{-\ln y}{1-y} \left(1 - \frac{4y(1+y)}{(1-y)^2} \right) + \frac{8y}{(1-y)^2} + \frac{1+y}{2} \right]. \quad (2.32)$$

When calculating the Compton scattering cross section using a computer, a problem called catastrophic cancellation arises. The function $f_C(y)$ features

two terms which go to infinity as the minimal energy fraction approaches unity. This is not a problem in real life, since the terms very nearly cancel each other. But when using floating point numbers, which have finite precision, to evaluate the expression, the subtraction leads to an error which is small compared with the individual terms, but large compared with the true value of the function. Since

$$y = \frac{m_e^2 c^4}{m_e^2 c^4 + 2Eh\nu(1 - \beta\mu)}, \quad (2.33)$$

we identify that the low-energy regime is the problem.

To correctly evaluate the function in all regions, we calculate the mathematical limit of the expression. By substituting $y = 1 - x$, we can calculate the limit as x approaches 0. Rewritten in terms of x , we have for the terms in the brackets

$$\frac{-\ln(1-x)}{x} \left(1 - \frac{4(1-x)(2-x)}{x^2} \right) + \frac{8(1-x)}{x^2} + \frac{2-x}{2}. \quad (2.34)$$

Using the Taylor expansion for the natural logarithm,

$$\ln(1-x) = -\sum_{n=1}^{\infty} \frac{x^n}{n}, \quad (2.35)$$

which holds for $-1 < x < 1$, we get

$$\frac{x + x^2/2 + x^3/3 + \mathcal{O}(x^4)}{x} \left(1 - \frac{4(1-x)(2-x)}{x^2} \right) + \frac{8(1-x)}{x^2} + \frac{2-x}{2}. \quad (2.36)$$

This simplifies to

$$1 + 2(1-x) - \frac{8}{3} + \frac{2-x}{2} + \mathcal{O}(x). \quad (2.37)$$

As x approaches zero, this approaches $4/3$. Hence, for y close to 1, we may use

$$f_C(y) \approx y \quad (2.38)$$

to avoid catastrophic cancellation. It should come as no surprise that Eq. (2.24) approaches the Thomson cross section in this limit.

2.2.4 Interaction Probability

In uniform background radiation, the distance travelled before an interaction occurs follows an exponential probability distribution with a mean free path of $1/\mathcal{R}$,

$$p(z) = \mathcal{R} \exp(-\mathcal{R}z), \quad (2.39)$$

where z is the distance travelled.

To simulate a cascade, we need a way to generate an interaction point for each particle. In the case of the exponential distribution, these can be generated

using the inverse transform sampling method. First, we calculate the cumulative probability distribution,

$$P(z) = \int_{-\infty}^z dx p(x) = 1 - \exp(-\mathcal{R}z). \quad (2.40)$$

Next, we assume that we can produce a random number r from a uniform distribution between 0 and 1. We convince ourselves, with the help of Fig. 2.8, that the probability that $P(z)$ is less than r is equal to r .

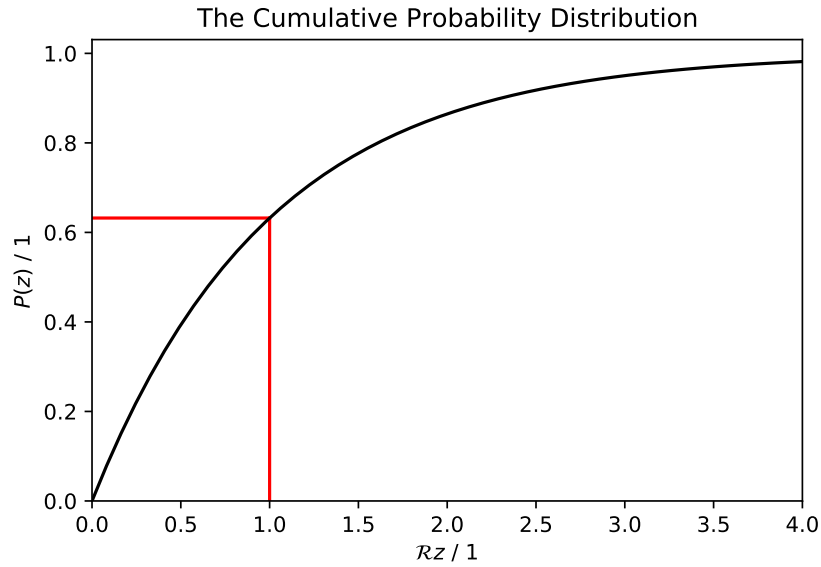


Figure 2.8: From the definition of the cumulative probability distribution, the probability that $\mathcal{R}z$ is less than 1 is equal to $P(1/\mathcal{R}) \approx 0.63$. The probability that the random number r between 0 and 1 is less than 0.63 is also 0.63. Hence, it clearly holds for any value of z that the probability that r is less than $P(z)$ is equal to $P(z)$. This means that, for random values of z sampled from $p(z)$, $P(z)$ and r are both uniformly distributed.

This means that we may generate a z -value by solving the following equation for z :

$$1 - \exp(-\mathcal{R}z) = r. \quad (2.41)$$

This amounts to inverting the function. We get that

$$z = -\frac{\ln(1-r)}{\mathcal{R}}. \quad (2.42)$$

We may simplify this slightly by recognizing that $1-r$ is also uniformly distributed between 0 and 1, which allows us to substitute it with r . The result

is

$$z = \frac{1}{\mathcal{R}} \ln \frac{1}{r}. \quad (2.43)$$

Inverse transform sampling is the method used in ELMAG. Since the background radiation is not uniform in our model, we need to find a different strategy. We introduce a finite step size Δz . At each step we calculate the probability that an interaction occurs between z and $z + \Delta z$. Should the interaction happen, we insert the resulting particles at $z + \Delta z$. This is illustrated in Fig. 2.9.




Figure 2.9: At each step we check if the random number r , between 0 and 1, is smaller than the interaction probability. If it is, an interaction occurs.

If the probability is too large, we lose information about the location where the interaction occurs. A step size of $4/\mathcal{R}$, for example, would mean that interactions occur within the interval with a probability of approximately 98% if the rate is fairly constant. The resulting interaction points, which would nearly all be located at the end of the first step, would do a very poor job of reproducing the correct probability distribution. We therefore require that the probability of interaction is small in each step.

2.3 The Accretion Disk

Black-body radiation from the accretion disk makes up the thermal background photons we will be considering. The temperature on the accretion disk can be modeled as

$$T = \left(\frac{GM\dot{M}}{4\pi\sigma R^3} \right)^{1/4}, \quad (2.44)$$

where G is the gravitational constant, M is the mass of the black hole, \dot{M} is the accretion rate, σ is the Stefan-Boltzmann constant, and R is the distance from the center [2, p. 36]. The temperature determines the spectrum at each point on the accretion disk through the Planck distribution, which ultimately leads to the interaction rates for pair production and inverse Compton scattering in Eqs. (2.22) and (2.31), respectively.

The characteristic size of a black hole is given by its Schwarzschild radius,

$$R_S = \frac{2GM}{c^2}. \quad (2.45)$$

A non-rotating black hole has an innermost stable circular orbit of $3R_S$ [6, p. 212]. This marks the inner edge of the accretion disk. Using Fig. 2.10, we may find the necessary parameters for evaluating the interaction rates through basic trigonometry.




Figure 2.10: The interaction site is located a distance z above the center of the black hole. The incident angle between the high-energy particle and the background radiation is θ . The background photon originates on the accretion disk, a distance R from the center of the black hole. The spectrum of background radiation depends on the temperature T at the point of origin. The inner radius R_i and outer radius R_o give rise to the minimal and maximal angles, θ_{\min} and θ_{\max} , respectively.

Azimuthal symmetry lets us express the angular limits solely in terms of the polar angle. The lower and upper limits of μ are respectively given by

$$\mu_{\min} = \cos \theta_{\max} = \frac{1}{\sqrt{1 + (R_o/z)^2}}, \quad (2.46)$$

where R_o is the outer radius, and

$$\mu_{\max} = \cos \theta_{\min} = \frac{1}{\sqrt{1 + (R_i/z)^2}}, \quad (2.47)$$

where R_i is the inner radius. The distance from the center of the accretion disk is given by

$$R = z \sqrt{\frac{1}{\mu^2} - 1}. \quad (2.48)$$

Chapter 3

The Program

The program is written in Fortran and can be found at github.com/vegahen in the repository `master_public`. It uses the Message Passing Interface (MPI), which allows for speeding up computation with parallel processes.

Unless otherwise specified, the figures in this chapter will be made for the values given in Tab. 3.1. The quantity M_{\odot} is the solar mass, which is 2×10^{33} g, and the Schwarzschild radius is calculated from the mass using Eq. (2.45).

Parameter	Value
R_i	$3R_S$
R_o	$100R_S$
M	$10^9 M_{\odot}$
\dot{M}	10^{27} g s^{-1}
R_S	$2.97 \times 10^{14} \text{ cm}$
z	10^{15} cm
E	10^{12} eV
μ	$\mu_{\max} \approx 0.75$

Table 3.1: AGN parameters and variables used in illustrations.

3.1 Integration

The first thing the program needs to do is calculate the interaction rate. For this purpose, a specialized algorithm for locating the integrand and determining appropriate step sizes for a numerical integration routine is constructed. Interaction rates are tabulated at the beginning of the program to be used during the propagation of particles.

The integration is carried out by the functions and subroutines contained in `integration101.f90`. This file contains two modules, `mu_integration` and `nu_integration`, in which the user parameters for the integration are defined

and the function values are stored. The definite integral up to each step is stored in order to construct the cumulative probability distributions for the interaction parameters. Function values are stored in a dynamically allocated array to ensure that the integration algorithm works with an arbitrary number of steps.

3.1.1 The Integrand

We will refer to both integrands as $g(\nu)$, and they consist of all terms given in Eqs. (2.22) and (2.31), including constants and multiplied with 2π since we have azimuthal symmetry. The integrands are calculated in the functions **integrand_pair** and **integrand_ics**. To avoid overflow, both check whether the exponent in the Planck distribution is larger than the logarithm of the largest double precision value, which is approximately $\ln(1.8 \times 10^{308}) \approx 709.8$.

Pair Production

It is determined whether we are above the pair production threshold. This is done by checking if

$$s = 2Eh\nu(1 - \mu) > 4m_e^2c^4, \quad (3.1)$$

and returning zero if it is not satisfied.

The function consists of a modified photon density, in which a factor of ν has been cancelled when dividing by s in the cross section, multiplied with $f_p(\beta_e)$, which is shown in Fig. 3.1.

Inverse Compton Scattering

The function consists of the photon density multiplied with $1 - \beta\mu$ and $f_C(y)$. The y -value above which the alternate expression in Eq. (2.38) is used is defined in the parameter **thomson_tol**. In Fig. 3.2 we observe the effect of using this substitution for $y > 0.999$.

3.1.2 Locating the Inner Integrand

The integrals with respect to ν are done in the functions **integrate_nu_pair** and **integrate_nu_ics**. The strategy used is roughly the same for both interactions, with minor differences to account for the properties of the cross sections.

Pair Production

We start by calculating the photon background frequency of the pair production cutoff. By solving $s = 4m_e^2c^4$ for ν , we find that this is given by

$$\nu_{\text{cutoff}} = \frac{2m_e^2c^4}{Eh(1 - \mu)}. \quad (3.2)$$




Figure 3.1: The terms in the brackets of the pair production cross section. We note that the function is zero below a certain frequency. This is what causes the pair production cutoff.

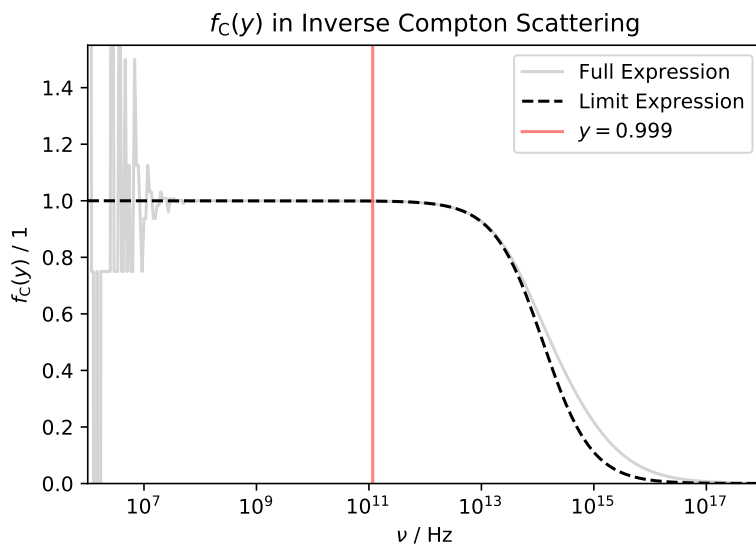


Figure 3.2: The dimensionless terms in the Compton scattering cross section. The gray curve shows catastrophic cancellation in $f_C(y)$ when using the full expression at lower frequencies.

The cutoff frequency is multiplied with the parameter **entrynu_scale_factor**, which is slightly smaller than 1, to ensure that the initial point we inspect is always below the pair production cutoff. The result is stored in the variable **nu_prev**.

Next we find a ν -value where $g(\nu) > 0$. Since the smallest possible floating point double precision value is $\epsilon = 4.9 \times 10^{-324}$, which we store in the parameter **epsilon**, we add some extra steps to ensure that we do not miss an integrand which has simply approached zero too quickly. The initial guess is twice that of ν_{prev} , and each guess follows the formula

$$\nu_{\text{curr}} = (1 + 2^{-i}) \nu_{\text{prev}}, \quad (3.3)$$

where i runs from 0 to the parameter **entrynu_searches**. Note that ν_{prev} does not change here. Once the first frequency to yield a nonzero function value is found, it remains stored in the variable **nu_curr**. If none of the guesses yield a nonzero value, the function is assumed to be zero.

Inverse Compton Scattering

For inverse Compton scattering, there is no cutoff at lower frequencies of the background radiation. In order to locate the function, we instead consider Wien's displacement law, which states that the frequency at which the spectral photon density has its peak is given by

$$\nu_{\text{Wien}} = 5.879 \times 10^{10} \text{ Hz K}^{-1} \cdot T. \quad (3.4)$$

Since $g(\nu)$ is the spectral photon density multiplied with $1 - \beta\mu$, which does not depend on ν , and $f_C(y)$, we can assume that the peak of the function is located near ν_{Wien} . The frequency is stored in the variable **nu_curr**. Looking at Fig. 3.2, we see that ν_{Wien} must always fall to the right of the peak, because $f_C(y)$ is strictly decreasing. Since we intend **nu_prev** to be on the left side of the peak, we make an initial guess of $\nu_{\text{Wien}}/10$.

3.1.3 Locating the Peak of the Inner Integrand

In order to determine a suitable step size for estimating the value of the inner integral, we locate the peak of $g(\nu)$. Once ν_{prev} and ν_{curr} have been determined, we add two more points between them such that the four are equidistant in ν . The ν -values are stored in the variables **b_left**, **b_midleft**, **b_midright**, and **b_right**, illustrated in Fig. 3.3. The function $g(\nu)$ has the property that there is only one maximum, something it inherits from the photon density. It is therefore possible to close in on the peak by looking at these four values of ν and their associated function values.

In order to start narrowing in on the peak, we need to ensure that $b_{\text{left}} < b_{\text{midleft}}$ and $b_{\text{right}} < b_{\text{midright}}$. If this is not the case, the algorithm incrementally moves the four points in the appropriate direction until the peak is enveloped. We then know that the peak must be located between b_{left} and b_{right} .

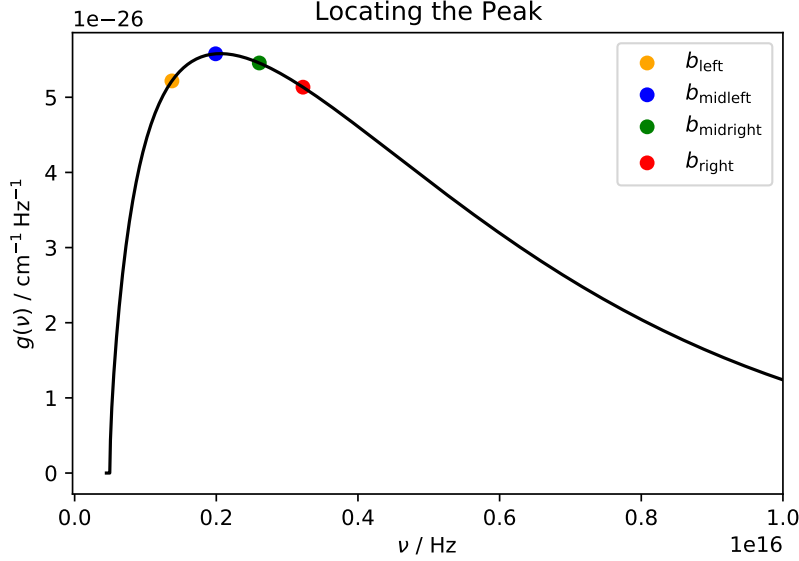


Figure 3.3: By evaluating the function at four different points and making sure that the peak is located between the outer points, we can always move closer to the peak by comparing the inner points.

We can approach the peak by comparing the function values of the middle points. If $g(b_{\text{midleft}}) < g(b_{\text{midright}})$, we know that the peak must be located to the right of b_{midleft} , hence we may set $b_{\text{left}} = b_{\text{midleft}}$. If the opposite is true, we set $b_{\text{right}} = b_{\text{midright}}$. The process is then repeated until we are arbitrarily close to the peak.

The parameters `peak_tol_nu_pair` and `peak_tol_nu_ics` specify how precisely we wish to locate the peak. Their values are compared with the variable `peak_proximity`, which is equal to

$$\frac{\min(g(b_{\text{left}}), g(b_{\text{right}}))}{\max(g(b_{\text{midleft}}), g(b_{\text{midright}}))}. \quad (3.5)$$

Once the proximity is greater than the specified tolerance, the variable `nu_b` is set equal to either b_{midleft} or b_{midright} , depending on which yields the larger function value.

3.1.4 Locating the Left Limit of the Inner Integrand

Pair Production

For pair production, the left limit of the integrand is located through a bisection technique. The variables `a_left` and `a_right` envelop the ν -value where the

integrand becomes zero, illustrated in Fig. 3.4. The function is then evaluated at $\nu_{\text{curr}} = (a_{\text{left}} + a_{\text{right}})/2$. If $g(\nu_{\text{curr}})$ is zero, we set $a_{\text{left}} = \nu_{\text{curr}}$, otherwise $a_{\text{right}} = \nu_{\text{curr}}$. The process is then repeated until we are arbitrarily close to the true pair production cutoff.




Figure 3.4: Bisection to find the left limit of the inner integral.

The parameter **foot_tol_nu_pair** specifies how precisely we wish to locate the cutoff. Its value is compared with the variable **peak_proximity**, which is equal to

$$\frac{g(a_{\text{right}})}{g(\nu_{\text{b}})}. \quad (3.6)$$

Once the proximity is smaller than the specified tolerance, the variable **nu_a** is set equal to a_{right} .

The program has the potential to get stuck once the function values get close to ϵ . If $\epsilon/g(a_{\text{right}})$ is greater than the tolerance, further iterations will not improve the proximity. In order for the program to continue, we stop bisection after a certain number of loops with $g(a_{\text{right}}) = \epsilon$, specified in the parameter **epsilon_loops**.

Inverse Compton Scattering

Since inverse Compton scattering has no cutoff, the left limit can be arbitrarily close to $\nu = 0$. The left limit is chosen by dividing a_{right} by ten until the variable **peak_proximity**, given by Eq. (3.6), is smaller than the parameter **foot_tol_nu_ics**. Afterwards, the variable **nu_a** is set equal to a_{right} .

3.1.5 Estimating the Inner Integral

We will refer to the inner integral as $h(\mu)$ in both interactions. With the left limit ν_a and the peak ν_b located, we may compute an initial estimate of the integral. This will be done using the trapezoid rule, which states that it is approximately equal to

$$h_{\text{est}}^T(\mu) = \sum_{i=1}^n \Delta h_i^T = \sum_{i=1}^n \frac{1}{2} (\nu_i - \nu_{i-1}) (g(\nu_{i-1}) + g(\nu_i)). \quad (3.7)$$

The general shape of $g(\nu)$ does not change as z , E , and μ are varied. We can therefore use ν_a and ν_b to determine appropriate step sizes $\Delta\nu_i = \nu_i - \nu_{i-1}$. The initial step sizes used are $\Delta\nu_1 = (\nu_b - \nu_a)/15$ for pair production and $\Delta\nu_1 = (\nu_b - \nu_a)/19$ for inverse Compton scattering. The initial step size is small to account for the sharp initial rise, and step sizes are doubled during the next two steps, after which it stays constant until $i = 8$ for pair production and $i = 9$ for inverse Compton scattering. For the remaining steps, the step sizes are $\Delta\nu_i = (\nu_i - \nu_a)/5$ for pair production and $\Delta\nu_i = (\nu_i - \nu_a)/6$ for inverse Compton scattering. The routine is illustrated in Fig. 3.5.




Figure 3.5: The initial step size is calculated based on $\nu_b - \nu_a$ and gradually increases.

The integrand decreases exponentially as ν becomes large. This lets us terminate when the average function value during the last step is smaller than some fraction of the average function value so far. This fraction is defined

in the parameters **estimate_tol_nu_pair** and **estimate_tol_nu_ics**. Stated mathematically, the criterium for terminating after step j is that

$$\frac{\Delta h_j^T / \Delta \nu_j}{\sum_{i=1}^j \Delta h_i^T / (\nu_j - \nu_a)}, \quad (3.8)$$

which is stored in the variable **estimate_proximity**, is smaller than the specified tolerance.


The trapezoid algorithm described above provides adequate precision for our present purposes, with an error of about 1%. The parameters **nmu_factor_pair** and **nmu_factor_ics** let us specify a smaller step size for the trapezoid rule if needed. When they are set equal to integer numbers larger than 1, each step in the original routine is divided into the specified number of steps. If they were for example set equal to 3, the steps 1, 2, and 3 would be of length $1/3 \cdot (\nu_b - \nu_a)/15$ for pair production and $1/3 \cdot (\nu_b - \nu_a)/19$ for inverse Compton scattering, while steps 4, 5, and 6 would be double of this. If it were of interest to compute the interaction rates more precisely, it would be better to implement a higher order method, e.g. Simpson's 3/8 rule. The step size could be determined in a similar way, initially as a fraction of $\nu_b - \nu_a$ followed by progressive increases, or it could be determined through an adaptive method.

3.1.6 Estimating the Outer Integral

The integrals with respect to μ are done in the functions **integrate_mu_pair** and **integrate_mu_ics**. They are much less complicated compared with those for ν , since the integral always runs between μ_{\min} and μ_{\max} , which are between 0 and 1. The estimate again uses the trapezoid rule, but this time with a constant step size and with the number of steps defined in the parameters **nmu_estimate_pair** and **nmu_estimate_ics**. The function is shown in Fig. 3.6.

Unlike for the inner integrand $g(\nu)$, the general shape of $h(\mu)$ varies for different values of the height z and the energy E and between the two different interactions. The maximum can be at either border of the interval or somewhere in the middle. At some heights and energies, the integrand is sharply peaked, so the error in the trapezoid estimate may become a problem. To check whether the number of steps is appropriate for the relevant heights and energies, one could compare the resulting interaction rates with those computed with a smaller step size, to see if the change is significant.

An attempt is made to increase the precision of the estimate slightly by determining the left and right limits where $h(\mu)$ goes to zero, stored in **p_left** and **p_right**. If the number of nonzero function values are less than the parameters **consecutivenonzeros_tol_pair** and **consecutivenonzeros_tol_ics**, the estimate is computed again with p_{left} and p_{right} as the integration limits.

Figure 3.6: The μ -integrand in the interaction rate.

3.1.7 Sampling the Angle and Frequency

Once it has been determined during propagation that an interaction between a photon or an electron with the background radiation happens, we need to determine the energy of the background photon and the relative angle between them. Since the location and width of the integrands on the frequency axis have a lot of variation, we will use the numerical version of the inverse transform sampling technique to determine the energy. For simplicity, this method is also used to sample the outer integral, but this leads to a considerable increase in runtime, since it involves many evaluations of the inner integral.

The unnormalized cumulative probability distributions are stored during integration in the arrays `mu_est` and `nu_est`. The subroutines `init_integration`, `delete_integration`, `increase_mu`, `increase_nu`, `add_mu`, and `add_nu` manage the dynamically allocated arrays in the integration modules. The integers `smu` and `snu` are the size of the dynamic arrays, while `nmu` and `nnu` are the number of elements currently stored. The arrays have dimension $s \times 2$, where s is the size, the first element of each subarray is the μ - or ν -value, and the second is the definite integral up to this step.

Sampling of μ and ν is done by interpolating the cumulative probability distributions from the integration. The integration functions take the float values `ran_mu` and `ran_nu` as input. If these are equal to -1 , no sampling is performed and the arrays are overwritten during the next integral without having been used. This is the case during tabulation of the interaction rates at the

beginning of the program. As particles interact, however, the integration functions are called with two random numbers between 0 and 1. These correspond to the percentiles of the cumulative probability distributions at which μ and ν are sampled, illustrated in Fig. 3.7. The resulting interaction parameters are returned through the parameters **mu_int** and **nu_int**.




Figure 3.7: With a random number r , we choose the μ -value where the area to the left is a fraction r of the total area below the curve. This sampling was performed in pair production for a photon with energy 10^{10} eV.

3.2 Propagation

The subroutine **propagate** repeatedly moves particles a small distance Δz along the z -axis, where it is ensured that the probability of interaction is less than a certain threshold. In addition to the constraint on the probability of interaction, there is a constraint on how much the interaction rate changes at each step. The particles are propagated either until they interact with the background photons, or until they escape to a specified distance **z_max** from the center of the black hole. The particles that initiate the cascades are injected at a height of **z_initial**.

3.2.1 Step Size

The initial step size is calculated using the interpolated interaction rate for the particle's height and energy. The interaction rate is stored in the variable **r1**.

Since the interaction rate is small, we may approximate it as

$$P(\Delta z) = 1 - \exp(-\mathcal{R}\Delta z) \approx \mathcal{R}\Delta z. \quad (3.9)$$

The probability which we consider sufficiently small is stored in the parameter **tol_prob**. Based on Eq. (3.9), we choose

$$\Delta z = 0.95 \cdot \frac{p_{\text{tol}}}{\mathcal{R}} \quad (3.10)$$

as the initial step size, with p_{tol} the specified tolerance. If the interaction rate is zero, it instead uses for Δz the parameter **incr_lots** multiplied with z . The height $z + \Delta z$ is stored in the variable **z_next**, and the interpolated interaction rate at this location is stored in the variable **r2**.

In addition to the restriction on the probability of interaction during a step, we also place a restriction on the change of the interaction rate from z to $z + \Delta z$. This change is stored in the variable **rise**, on the form

$$\left| \ln \frac{\mathcal{R}_1}{\mathcal{R}_2} \right|, \quad (3.11)$$

and is assured to be less than the parameter **tol_rise**. Should the probability of interaction drop below the parameter **tol_prob_highrise**, this criterium is ignored to avoid slowing down the propagation. Similarly, if either \mathcal{R}_1 or \mathcal{R}_2 is zero, the step is accepted if the probability is less than the parameter **tol_prob_zero**.

Since we have a restriction on the change in the interaction rate, using Eq. (3.10) at each step would fail if the rate changes too rapidly. Instead, we will use an adaptive routine where the change in the step size depends on the outcome in the previous iteration. If the previous step failed, the step size is reduced by a factor stored in the parameter **decr** and another attempt is made. Should the probability of interaction be zero in the previous step, it will use for Δz , as it did initially, the parameter **incr_lots** multiplied with the current height z . If the probability was less than **tol_prob_verylow**, the step size is increased by the parameter **incr_more**. Otherwise, it is increased by the parameter **incr**.

3.2.2 Interactions

For a given step, the probability of interaction used is $(\mathcal{R}_1 + \mathcal{R}_2)/2 \cdot \Delta z$. A random number is generated for each accepted step, and an interaction occurs if it is smaller than this probability.

Interactions are handled by the subroutine **interaction_agm**. First, two random numbers are generated, followed by a call to either **integrate_mu_pair** or **integrate_mu_ics** to sample the interaction parameters μ and ν . Next, the energy fractions taken by the resulting particles are sampled. This is done by the functions **zpair** and **zics**, which are taken directly from ELMAG. The distributions for the energy fractions are shown for several energies in Figs. 3.8

and 3.9. The sampling was done after picking the most likely value of μ followed by picking the most likely value of ν at the specified energies, with 10 million samples for each curve.

As seen in Fig. 3.9, interactions of lower-energy electrons will result in soft photons, which are photons scattered with energies much lower than that of the initial electron. This poses a problem to the simulation because very many interactions will happen without significantly changing the energy of the electron. In ELMAG this is solved by computing a modified interaction rate which contains a maximum energy fraction y_{\max} . This eliminates the contribution from interactions that yield low energy photons, and the interaction rate goes to zero somewhere in the Thomson regime. Instead, we specify the two parameters **electron_energythr** and **photon_energythr**. Whenever an interaction produces a particle with an energy below threshold, the particle is discarded instead of being tracked further. When looking at the resulting photon spectrum, we should therefore keep in mind that for photon energies below $E_{\text{thr},e}(1-y)$, where y is calculated for the electron energy threshold, the spectrum is no longer accurate.

3.2.3 Storing Particles

All data associated with a particle is stored in the custom data type **event**, which is equivalent to the **one_event** data type in ELMAG. The integer **icq** contains the particle type, with 0 for photons, 1 for electrons, and -1 for positrons. Each particle has a unique value of the integer **id**, which is given in ascending order as particles are created, starting with the number 1. The variables z and e contain the height above the accretion disk and the energy of the particle, respectively, while w is a weight given to each particle.

The weight w can be used to implement weighted sampling, something that was done in ELMAG. It is also used during injection, to let us use a power law spectrum of particles to initiate cascades. The energy of injected particles is determined in the subroutine **inject**. Since we are interested in how initial energies across several decades affect the interactions of the cascade, these are sampled uniformly on a logarithmic scale. Uniform logarithmic sampling corresponds to a power law spectrum of $dN/dE = E^{-1}$. In order to obtain an arbitrary $E^{-\alpha}$ power law spectrum from this, each particle is given a weight which is proportional to $E^{-(\alpha-1)}$.

Once particles are created, they are stored in a dynamic heap structure. It is a min heap which arranges particles based on their energies, which means that the first element in the dynamic list of particles is always the particle with the smallest energy. This is done to limit the number of particles that need to be stored at any given time, by always propagating the particle with the lowest energy first to avoid a buildup of low-energy particles in the heap.

The heap itself is located in the module **heap_agm**. The variable **s** is the size of the dynamic array of **event** structures, while **n** is the number currently stored. The file **heap101.f90** contains both the subroutines that handle the dynamic array and the subroutines that handle the insertion and removal of elements.

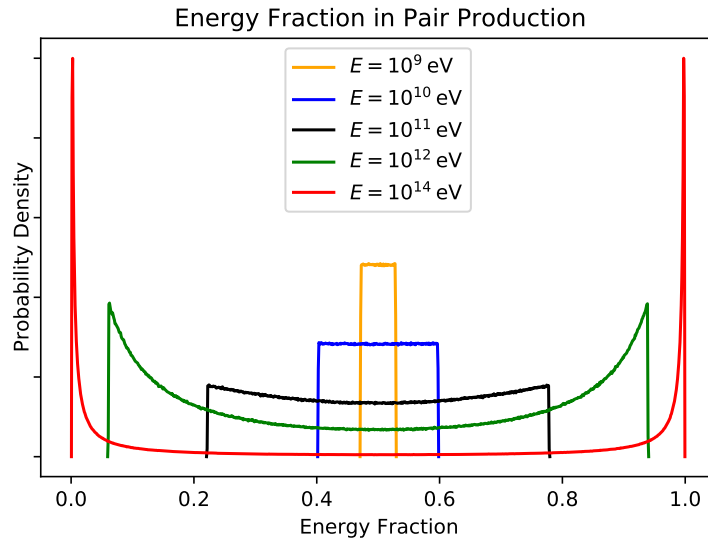


Figure 3.8: The probability distributions for the energy fraction taken by one of the electrons in pair production. The probability densities are not in proportion with one another.

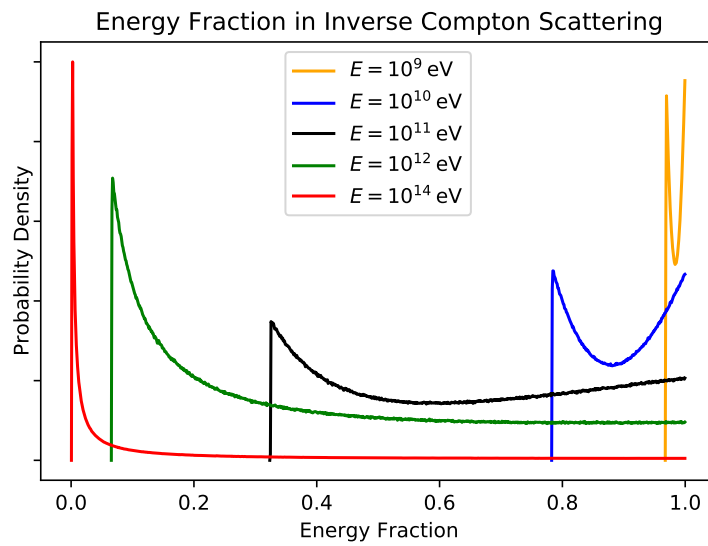


Figure 3.9: The probability distributions for the energy fraction taken by the electron in inverse Compton scattering. The probability densities are not in proportion with one another.

The module also contains integer variables for the number of particles injected, the number of particles created (the number which at any time represents the **id** value of the most recently created particle), the number of particles that have escaped, and the number of particles that have been ignored due to having energies below threshold.

Once a particle escapes to the specified distance **z_max**, the energy and weight of the particle is added to the resulting spectrum in the subroutine **store**. The weights of the particles are added to energy bins which are of equal size on the logarithmic scale. The spectrum is stored on the form $E^2 dN/dE$. To achieve this with energy bins which have a width already proportional to E , we only need to multiply with one additional factor of E .

3.3 Program structure

User variables are found in the file **modules101.f90**, with the exception of integration parameters, which are found in **integration101.f90**, and the injection energies, which can be changed directly in the **inject** subroutine in **init101.f90**. The primary user variables are found in the module **user_variables**. The parameter **n_start** specifies the number of particles which are injected in each set, while **n_sets** is the number of sets. At the start of the program, the subroutine **init** in **init101.f90** is called to initialize the necessary variables for the simulation. The seed for the pseudo-random number generator, **iseed**, is different for each parallel process when using MPI.

Interaction rates tabulated for interpolation are calculated at the start of the program. This is done on a grid of heights and energies which are both spread uniformly on the logarithmic scale. They are stored in the files specified by **agn_pair_filename** and **agn_ics_filename**, and in subsequent runs they can be read from these files by letting the parameter **readagnfit** be true. The filenames, as well as the following grid parameters, are found in the module **agn_fit**. The parameters **zmin_pair** and **emin_pair** specify the logarithm of the minimal heights and energies which the rates are calculated for in pair production, while **d_zpair** and **d_epair** specify the number of gridpoints per decade, and **n_zpair** and **n_epair** specify the total number of gridpoints. For example, a minimum energy of 8 with $d = 10$ and $n = 31$ would produce a grid with energies ranging from 10^8 eV to 10^{11} eV. Equivalent parameters are used for the inverse Compton scattering grid. The heights, energies and rates are stored as logarithms to simplify the interpolation subroutines, which are found in **interpolation101.f90**, since the interpolation is done linearly on the logarithmic scale.

The modules **internal** and **result** contain the variables used to create the spectrum of escaping particles. The minimal energy on the spectrum is given in **e_min**, while **dn** is the width of each energy bin on the logarithmic scale (for example, 0.1 means that there are 10 bins per decade), and **n_enbin** is the total number of energy bins. If MPI is used, the results for different processes are added to the array **en_f_tot** at the end of each set. The spectra are written

to text files in **output101.f90**, and the files are stored in the "Data" folder in the project directory. We are only interested in the photon spectrum, which has "gam" in its name. The first and second columns are the primary and secondary axes of the spectrum, the primary in units of electronvolts, and the secondary in arbitrary units. The third and fourth columns are the logarithms of the first and second.

Mass, accretion rate, and accretion disk size are found in the **agn_data** module. Constants are stored in the module **constants**. The parameters for the propagation subroutine are found in the module **propagation_params**.

The module **test_module** is used for debugging the program. The parameter **show_interactions** prints all particles as they are removed from, added to, or ignored by the heap. The parameter **show_propagation_wr** prints the parameters of each interaction. The parameter **show_propagation_i** makes the program interactive by waiting after each interaction. The user inputs a number which specifies the next interaction to stop at, but it can be turned off by entering a number smaller than 1. The parameter **show_iterations** lets us interactively look through each step in the propagation. Entering a number smaller than 1 turns off this functionality. The parameter **show_integration** prints all values added to the integration lists.

Chapter 4

Results and Discussion

The escaping photon spectrum for the parameters specified in Tab. 3.1 is shown in Fig. 4.1. It is the result of injecting 10 photons at $3R_S = 8.91 \times 10^{14}$ cm with energy 10^{13} eV using an electron energy threshold of 10^7 eV. 43 million particles are created.

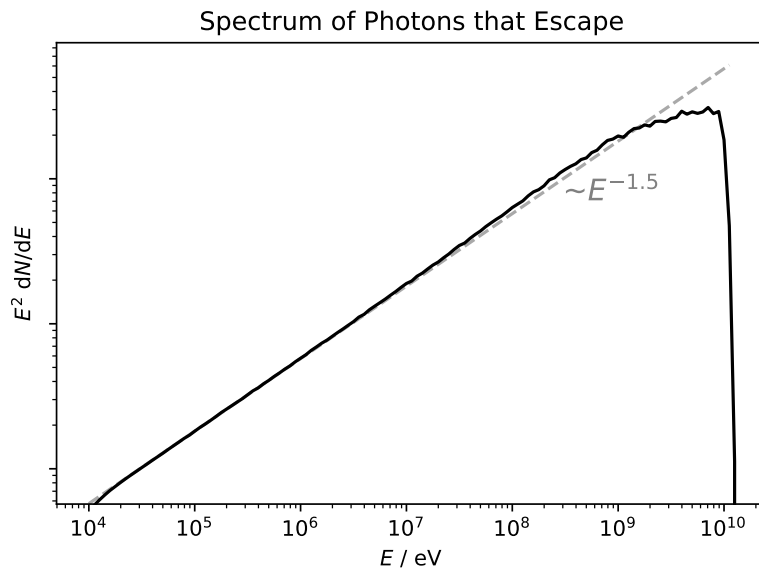


Figure 4.1: The spectrum that result from cascades started by 10 photons with energy 10^{13} eV.

The theoretical prediction for a uniform background is that there should be a power law spectrum with $\alpha = -1.5$ at lower energies [1, p. 2]. The photon spectrum deviates slightly from the $E^{-1.5}$ power law between 10^7 eV and 10^9 eV.

By analyzing the interactions at an electron energy of 10^7 eV, we find that the probability distribution for ν is shifted to higher energies as μ increases. At $\mu = \mu_{\max}$, the 99th percentile of the frequency of the background photons which interact with electrons is 2.25×10^{16} Hz, suggesting that this is close to the maximal background frequency contributing to interactions. The minimal energy fraction found using this frequency, from Eq. (2.33), is $y = 0.998$, suggesting that photons with an energy below approximately 2×10^4 eV could still be produced from discarded electrons. Looking at the spectrum in Fig. 4.1, the curve bends slightly in this area, supporting the idea that photons are missing. Decreasing the electron energy threshold does indeed straighten the curve in this region.

Particles are propagated to $z_{\max} = 10^{16}$ cm, but most of the interactions happen near 10^{15} cm. To see this we could simply change the value of z_{\max} and compare the results, or we could store the z -values at each interaction, the result of which is shown in Fig. 4.2.

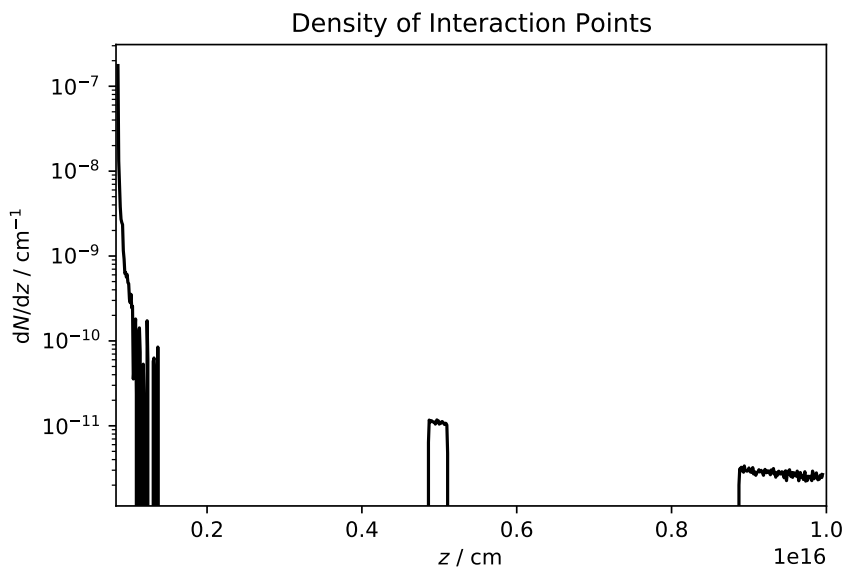


Figure 4.2: The distribution of interaction points, with two nonphysical trailing cascades.

The cutoff in pair production happens around 7×10^9 eV in the simulation above. The two trailing sub-cascades are initiated by photons with energies of about 6×10^9 eV, but these interactions are caused by step sizes that are too large, and therefore overestimate the interaction probability. This happens because the algorithm ignores the criterium on the change in the interaction rate if the probability is very small, which is the reason it only happened twice. We could adjust the propagation parameters to reduce this effect, but always

insisting that the rate cannot change much will drastically increase the number of steps.

The pair production rate near cutoff as a function of the height, seen in Fig. 4.3, shows that it is virtually impossible that the two photons interact where they do. The reason they appear in this region is that interactions happen at the end of the propagation step, while the probability is calculated using the average of R_1 and R_2 . From the inverse Compton scattering rate in Fig. 4.4, we also see that the scattering rate is much smaller in the second trailing cascade than in the first. This explains why the Compton scattering interactions in the second cascade happen further apart, as we can see from the smaller density of interaction points.

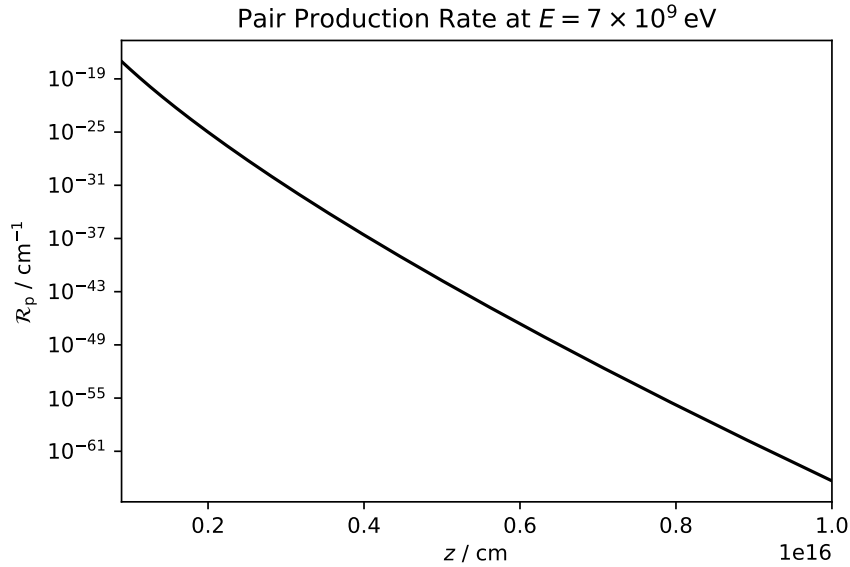


Figure 4.3: The pair production rate as a function of the height near the pair production cutoff energy.

We may compare the pair production cutoff with that of the project thesis, which was on the topic of optical depth in the same AGN system. The optical depth was calculated from the center of the system, so if we inject particles at smaller z -values, we should get similar results. The cutoff in pair production with $z_{\text{initial}} = 0.03R_S = 8.91 \times 10^{12}$ cm happens at around 3×10^9 eV. The optical depth was found to be of order 1 near 4×10^9 eV in the project.

Although the interaction rates close to the pair production cutoff energy change very rapidly as z increases, the rate at other energies can be fairly constant across the typical interaction length. A more effective propagation algorithm in these regions could be to locally use inverse transform sampling, instead of using many propagation steps to traverse a region with nearly constant

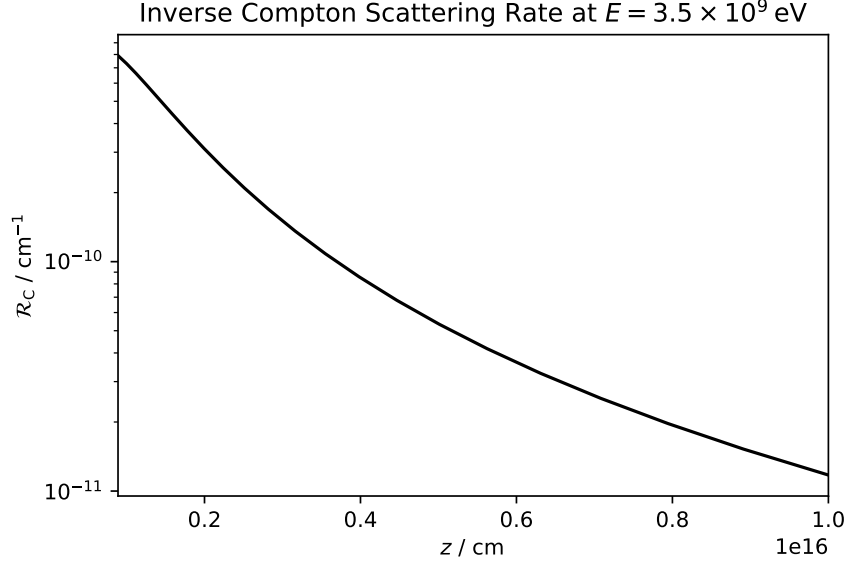


Figure 4.4: The inverse Compton scattering rate as a function of the height for electrons produced close to the pair production cutoff.

rates. This would allow for much larger step sizes that are only limited by the change in the interaction rate, making it much more effective when the rate is so large that step sizes determined by the probability tolerance, roughly given by Eq. (3.10), are much smaller than the length over which the rate changes considerably.

The most time consuming part of the program is currently interactions, since it calculates the rate to sample μ . The simulation above created more than 40 million particles, but the vast majority of these are low energy photons produced in the Thomson scattering regime. By choosing to not keep some of the electrons with lower energies and increasing the weight of the electrons we do keep, the number of interactions could be decreased drastically. This is weighted sampling, which has been implemented in ELMAG. Yet another way to speed up calculations would be to construct a more effective way to sample μ , such as by storing the probability distributions during tabulation and interpolating them later, or by using a more effective Monte Carlo method.

To obtain a smooth and detailed spectrum from a Monte Carlo simulation, we need a large number of photons created in each region. Although weighted sampling has not been implemented, we can obtain a more detailed spectrum manually by combining the spectra made with different electron energy thresholds. This way, we can ensure that many photons are created in all the different energy regimes. The result is shown in Fig. 4.5. The general relationship found between the threshold and the number of particles is that they are inversely

proportional to each other, so if the electron energy threshold is reduced by a factor of ten, there will be approximately ten times as many particles created.




Figure 4.5: A spectrum obtained by combining simulations with different electron energy thresholds. As the threshold increases, we increase the number of injected particles to ensure that enough photons are created in each region. A total of 2 billion particles were created to produce this plot.

As seen in Fig. 4.3, the program is able to calculate interaction rates that are incredibly small. A simpler method for calculating the rate used in the project thesis, where step sizes do not depend on heights nor energies, is compared with the new method in Fig. 4.6, as functions of energy. It is Eq. (3.3) that can take the credit for finding the smaller integrands. The step size in the simple method is so large that at lower energies it misses the integrand completely. In fact, the new method is able to calculate pair production rates down to approximately $10^{-275} \text{ cm}^{-1}$, but this is of course not useful in any practical sense.

The inverse Compton scattering rate as a function of energy is shown in Fig. 4.7. We note that the interaction rate is constant in the Thomson scattering regime, but as the energy approaches the rest mass of the electron, it starts to increase.

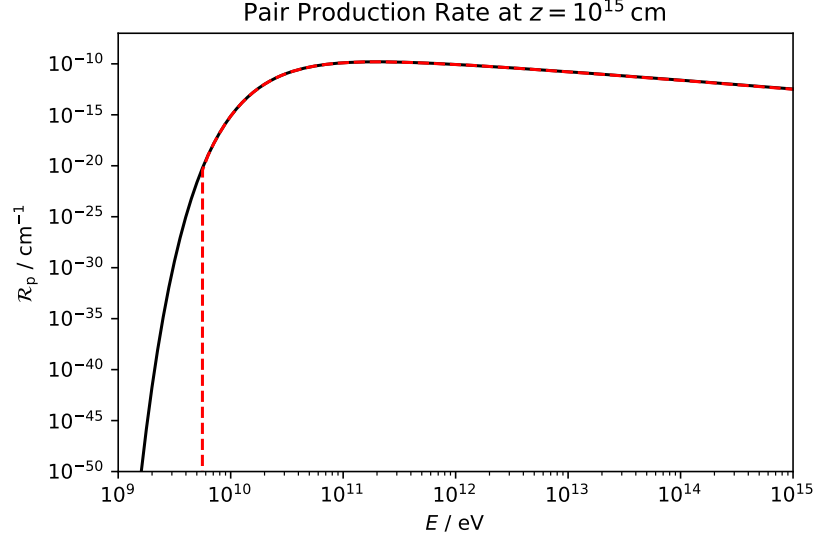


Figure 4.6: The pair production rate computed using a traditional integration routine, in red, versus computed using the program described in this thesis.




Figure 4.7: The inverse Compton scattering rate computed with the program described in this thesis.

Chapter 5

Conclusion

The program is designed to be a robust algorithm which solves the problem presented. Considerable effort has been made to ensure that interaction rates can be computed accurately for nearly any parameters and that the algorithm works for any type of changing interaction rates. It should be possible to expand the program to work for arbitrary backgrounds without too much effort, easily so if there is azimuthal symmetry. There is considerable potential for improving the runtime, but the initial version can serve as a baseline to check that the methods used to speed up calculations are working as they should.

Weighted sampling could be implemented to remove many of the particles at lower energies. Since every interaction results in two particles, the number of low-energy particles grows extremely large as the cascade progresses. By only keeping a certain fraction of representative low-energy particles and modifying their weight to make up for those lost, the runtime could be drastically improved.

The μ -value could be sampled in a more effective way. The current program computes the rate at the interaction point in order to construct the probability distribution, which is very computationally intensive and the current bottleneck of the entire program. One possibility is to store the probability distributions along with the tabulated interaction rates and interpolate them later. Another is to use an altogether different Monte Carlo method.

Instead of a purely stepwise propagation of particles, one could implement a hybrid approach which uses inverse transform sampling to find interaction points on an interval where the rate is almost constant. In many situations the interaction rate does not change much during a typical interaction length, which means that locally we could consider it to be uniform.

Bibliography

- [1] M. Kachelrieß, S. Ostapchenko and R. Tomàs, *Elmag: A monte carlo simulation of electromagnetic cascades on the extragalactic background light and in magnetic fields*, *Computer Physics Communications* **183** (2012) 1036.
- [2] B.M. Peterson, *An Introduction to Active Galactic Nuclei*, Cambridge University Press (1997).
- [3] G.A. Shields, *A Brief History of Active Galactic Nuclei*, *PASP* **111** (1999) 661.
- [4] P.C. Hemmer, *Termisk fysikk*, Vigmostad & Bjørke AS, 3rd ed. (1989).
- [5] G. Sigl, *Astroparticle Physics: Theory and Phenomenology*, Atlantic Press (2017).
- [6] S. Carroll, *Spacetime and Geometry: An Introduction to General Relativity*, Pearson, new international ed. (2014).

