Ultra-high-energy cosmic-ray nuclei and neutrinos in models of gamma-ray bursts and extragalactic propagation

DISSERTATION

zur Erlangung des akademischen Grades
doctor rerum naturalium
(Dr. rer. nat.)
im Fach: Physik
Spezialisierung: Theoretische Physik

eingereicht an der
Mathematisch-Naturwissenschaftlichen Fakultät
der Humboldt-Universität zu Berlin

von

M.Sc. Jonas Heinze

Präsidentin der Humboldt-Universität zu Berlin:
Prof. Dr.-Ing. Dr. Sabine Kunst

Dekan der Mathematisch-Naturwissenschaftlichen Fakultät:
Prof. Dr. Elmar Kulke

Gutachter: 1. PD Dr. Walter Winter
2. Prof. Dr. Thomas Lohse
3. Prof. Dr. Günter Sigl

Tag der mündlichen Prüfung: 27. Januar 2020

Selbständigkeitserklärung

Ich erkläre, dass ich die Dissertation selbständig und nur unter Verwendung der von mir gemäß § 7 Abs. 3 der Promotionsordnung der Mathematisch-Naturwissenschaftlichen Fakultät, veröffentlicht im Amtlichen Mitteilungsblatt der Humboldt-Universität zu Berlin Nr. 42/2018 am 11.07.2018, angegebenen Hilfsmittel angefertigt habe.

Abstract

Ultra-high-energy cosmic rays (UHECRs) are the most energetic particles observed in the Universe. Their energy spectrum, chemical composition and arrival directions are observed in extensive air-shower (EAS) experiments, notably the Pierre Auger Observatory (Auger) and Telescope Array (TA). While the astrophysical sources of UHECRs have not yet been uniquely identified, there are strong indications for an extragalactic origin. The interpretation of the observations requires both simulations of UHECR acceleration and energy losses inside the source environment as well as interactions during extragalactic propagation. Due to their extreme energies, UHECR will interact with photons in these environments, producing a flux of secondary neutrinos. The IceCube observatory has detected a neutrino flux of astrophysical origin, which likely consists of neutrinos from astrophysical sources and not of neutrinos produced during extragalactic propagation of UHECRs (so-called 'cosmogenic' neutrinos).

This dissertation deals with models of UHECR sources and the accompanying neutrino production in the source environment and during extragalactic propagation. At present, there is a wide range of possible UHECR source candidates with different UHECR acceleration and interaction mechanisms. Therefore, the generic source properties imposed by propagation are studied in this thesis, as well as gamma-ray burst (GRBs) as one specific source candidate. These models depend on the cross sections for photo-nuclear and hadronic interactions of UHECRs, which are only sparsely measured. This affects the photo-disintegration and photo-meson production both in the source and during extragalactic propagation. Furthermore, the EAS development is driven by hadronic interactions of UHECRs with air atoms, which introduces degeneracies in the interpretation of the observed data in terms of chemical composition.

We have developed a new, computationally efficient code, PriNCe, for the extragalactic propagation of UHECR nuclei. The computational speed allows to efficiently vary the cross sections and photon-environments between runs. The PriNCe code is applied for an extensive parameter scan of a generic source model that is described by the spectral index, the maximal rigidity, the cosmological source evolution and the injected mass composition. In this scan, we demonstrate the impact of different disintegration and air-shower models on the inferred source properties. A prediction for the expected flux of cosmogenic neutrinos is derived, which is found to be out of range for future neutrino detectors, independent of the model uncertainties.

GRBs are discussed as specific UHECR source candidates in the multi-collision internal-shock model. This model takes the radiation from different radii in the GRB outflow into account. We demonstrate how different assumptions about the initial setup of the jet and the hydrodynamic collision model impact the production of UHECRs and neutrinos. Motivated by the multi-messenger observation of GRB170817A, we discuss the expected neutrino production from this GRB and its dependence on the observation angle. We show that the neutrino flux for this event is at least four orders of magnitude below the detection limit for different geometries of the plasma jet.

Kurzzusammenfassung

Utrahochenergetische kosmische Strahlung (ultra-high-energy cosmic rays – UHECR) besteht aus ionisierten Atomkernen mit den höchsten Teilchenenergien, die je gemessen wurden. Sie erzeugen Luftschauer in der Erdatmosphäre, anhand derer Experimente wie das Pierre Auger Observatory (Auger) und das Telescope Array (TA) ihre Energien, chemische Komposition und Ankunftsrichtungen rekonstruieren. Zwar wurden die Quellen von UHECRs noch nicht eindeutig identifiziert, doch gibt es deutliche Anzeichen, dass sie extragalaktisch sind. Um die Beobachtungen zu interpretieren, wird sowohl ein Modell der UHECR Beschleunigung und Wechselwirkungen in der Quelle als auch der Wechselwirkungen während der extragalaktischen Propagation benötigt. Dabei werden astrophysikalische Neutrinos erzeugt, welche vom IceCube Observatory nachgewiesen wurden. Der gemessene Fluss besteht wahrscheinlich aus Neutrinos aus astrophysikalischen Quellen und nicht aus sogenannten kosmogenischen Neutrinos, die während der UHECR Propagation entstehen.

Diese Dissertation behandelt Modelle der Quellen von UHECRs und die damit verbundene Produktion von Neutrinos sowohl in den Quellen als auch während der Propagation. Momentan gibt es eine Vielzahl von möglichen Quellen von UHECRs, in denen verschiedene Mechanismen zur Beschleunigung und Wechselwirkung von UHECRs führen können. In dieser Arbeit werden daher sowohl die generellen Eigenschaften der Quellen anhand der Propagation abgeleitet als auch Gammastrahlenblitze (gamma-ray bursts – GRBs) als eine spezielle Quelle behandelt. Für photohadronische Wechselwirkungen, welche einen wichtigen Einfluss auf die Disintegration in der Quelle und während der Propagation haben, liegen nur unzureichende Messungen vor. Zusätzlich hängt die Interpretation der chemischen Komposition in Luftschauerexperimenten von Modellen der hadronischen Wechselwirkungen von UHECR Atomen mit der Luft ab.

Für diese Arbeit wurde ein neuer Code, PRINCE, für die Propagation von UHECRs entwickelt, der ermöglicht, die Wirkungsquerschnitte und Photonumgebungen effizient zwischen einzelnen Berechnungen zu variieren. Dieser Code wird in einem umfangreichen Parameterscan für ein generisches Quellenmodell angewendet, welches mit dem Spektralindex, der maximalen Rigidität, der kosmologischen Quellenverteilung und der chemischen Komposition als freie Parameter definiert ist. Dabei wird der Einfluss von verschiedenen Photodisintegrations- und Luftschauermodellen auf die erwarteten Eigenschaften der Quellen demonstriert. Der Fluss kosmogenischer Neutrinos, welcher sich aus dem erlaubten Parameterraum ableiten lässt, liegt unabhängig von den Modellunsicherheiten außerhalb der Reichweite aller derzeit geplanten Neutrinodetektoren.

GRBs als mögliche Quellen von UHECRs werden im Multi-Collision Internal-Shock Modell simuliert, welches die Abhängigkeit der Strahlungsprozesse von den verschiedenen Dissipationsradien im Plasmajet berücksichtigt. Für dieses Modell wird der Effekt demonstriert, den verschiedene Annahmen über die anfängliche Verteilung des Plasmajets und das hydrodynamische Modell auf die resultierende UHECR- und Neutrinosstrahlung haben. Für den Gammastrahlenblitz GRB170817A, welcher zusammen mit einem Gravitationswellensignal beobachtet wurde, werden Vorhersagen für den Neutrinofluss und ihre Abhängigkeit vom Beobachtungswinkel gemacht. Diese liegen mindestens vier Größenordnungen unterhalb der Detektionslimits.

Acknowledgements

First, I would like to thank Walter Winter for his supervision and support during the three years of this PhD project. He taugh me alot about how to work as a scientist as well as about multi-messenger astrophysics in particular. I want to thank everyone in the NEUCOS group that I worked with and learned from. A special thanks goes to Anatoli Fedynitch and Denise Boncioli who I worked most closely with and who both taught me many different things about science. I also appreciate the collaboration with all of my co-authors in the publications written during this project, who have not been mentioned yet: Daniel Biehl, Mauricio Bustamante, Annika Rudolph, and Kohta Murase.

I wish to thank everyone who helped me during the writing of this thesis. Denise, Anatoli, Arjen, Shan and Annika for reading different parts of the draft versions and providing valuable feedback. Thanks also to Sasha and Christina for proofreading the almost final version.

Finally, I want to thank all of my friends and family, who have given me support and distraction during the sometimes exhausting times of this project. I want to thank both my parents who have always supported me. Without them I would surely not have made it to this point.

During my work leading to this thesis, I was supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (Grant No. 646623)

List of publications

The following articles were published in connection with this dissertation.

Peer-reviewed publications:

- Mauricio Bustamante, Kohta Murase, Walter Winter, and Jonas Heinze. Multi-messenger light curves from gamma-ray bursts in the internal shock model. *Astrophys. J.*, 837(1):33, 2017. doi: 10.3847/1538-4357/837/1/33
- Daniel Biehl, Jonas Heinze, and Walter Winter. Expected neutrino fluence from short Gamma-Ray Burst 170817A and off-axis angle constraints. *Mon. Not. Roy. Astron. Soc.*, 476(1):1191–1197, 2018. doi: 10.1093/mnras/sty285
- Jonas Heinze, Anatoli Fedynitch, Denise Boncioli, and Walter Winter. A new view on Auger data and cosmogenic neutrinos in light of different nuclear disintegration and air-shower models. *Astrophys. J.*, 873(1):88, 2019. doi: 10.3847/1538-4357/ab05ce
- Annika Rudolph, Jonas Heinze, Anatoli Fedynitch, and Walter Winter. Impact of the collision model on the multi-messenger emission from Gamma-Ray Burst internal shocks. *Astrophys. J.*, 893(1):72, 2020. doi: 10.3847/1538-4357/ab7ea7

The following article was published before the start of this PhD-project

- Jonas Heinze, Denise Boncioli, Mauricio Bustamante, and Walter Winter. Cosmogenic Neutrinos Challenge the Cosmic Ray Proton Dip Model. *Astrophys. J.*, 825(2):122, 2016. doi: 10.3847/0004-637X/825/2/122

List of abbreviations

AGN	active galactic nucleus
CIB	cosmic infrared background
CMB	cosmic microwave background
CR	cosmic ray
DOM	digital optical module
EAS	extensive air shower
EBL	extragalactic background light
GRB	gamma-ray burst
GW	gravitational wave
GZK	Greisen-Zatsepin-Kuzmin
HESE	high-energy starting event
SED	spectral energy distribution
SNR	supernova remnant
SBG	starburst galaxy
TDE	tidal disruption event
UHE	ultra-high-energy
UHECR	ultra-high-energy cosmic ray

Contents

Chapter 1

Introduction

Cosmic rays are extraterrestrial charged particles, mainly protons and nuclei, that constantly hit the Earth's atmosphere. Their observed energies range over many orders of magnitude from below 1 GeV to the most energetic particles ever observed at 10^{11} GeV. The steeply falling flux constitutes a challenge for detection at higher energies. At energies above $\sim 10^5$ GeV, cosmic rays can only be observed indirectly through the air showers they initiate in the atmosphere. Cosmic rays above $\sim 10^8$ GeV are typically referred to as UHECRs. Their energy, mass composition and arrival direction are currently measured by the Pierre Auger Observatory [6–8] and the Telescope Array [9–11], which are extensive air-shower detectors.

The apparent isotropy in the arrival directions of UHECRs is strong evidence that they are of extragalactic origin [12, 13]. This is complemented by the fact that no known source in our galaxy is capable of accelerating UHECRs to their extreme energies. The identification of the sources of UHECRs is complicated because UHECRs are deflected in extragalactic and Galactic magnetic fields, losing the direct spatial correlation to their sources [14]. While the sources have therefore not yet been uniquely identified, there are several viable source candidates that are in principle luminous and abundant enough to power the flux of UHECRs [15, 16]. In most of these source environments as well as during extragalactic propagation, UHECRs will interact with photon fields, leading to energy losses and the production of secondary neutrinos and gamma-rays. For UHECR nuclei this also leads to photo-disintegration.

The mass composition measurements by the Auger collaboration show a significant abundance of UHECR nuclei, with the average composition getting heavier at the highest energies [17]. To interpret these observations in terms of the source physics, a complete model of the diffuse flux and mass composition of UHECRs needs to take into account both the acceleration and interactions inside the sources as well as the interactions during extragalactic propagation. Since the disintegration cross sections are only measured for a few stable elements, this adds significant uncertainties to the model [18]. Additionally, the reconstruction of mass composition from the air-shower profile relies on hadronic interaction models [19, 20]. These models add additional uncertainties, as the first few interactions in the air shower happen at center-of-mass energies that are several orders of magnitude above those obtainable at the LHC. The largest uncertainty, however, comes from the model of the source environment, since the different possible source candidates can have widely different properties like size, radiation density and distribution over redshift.

Taking into account additional messengers can help to constrain some of these uncertainties. Specifically neutrinos are a signature of UHECRs at the source, since they are produced dominantly in photo-hadronic interactions. A flux of neutrinos with astrophysical origin has in fact been observed by the IceCube Observatory [21, 22]. As neutrinos are electrically neutral, they are not deflected by magnetic fields and can point towards the sources of UHECRs. While no clear correlation has been established yet, the non-observation of neutrinos from specific source classes leads to stacking limits [23–25]. Similarly, the non-observation of neutrinos produced during extragalactic propagation, so-called ‘cosmogenic’ neutrinos, limits both the cosmological evolution of sources and the UHECR composition. This is because cosmogenic neutrinos reach us even from high redshifts, while UHECRs do not, and are produced much more copiously from UHECR protons than from heavier nuclei [5, 26, 27].

In this thesis, we discuss both the extragalactic propagation of UHECRs as well as gamma-ray bursts (GRBs) as possible sources of UHECRs. For this, we developed a new code, PriNCe, to simulate the extragalactic propagation of UHECRs and the production of cosmogenic neutrinos. This code is applied in a comprehensive parameter scan of the main model uncertainties affecting the interpretation of UHECR data. GRBs are discussed in a sophisticated multi-collision model, which takes into account collisions in the plasma jet at different radii. From this model we derive the expected neutrino and UHECR output.

The thesis is structured as follows: In Chapter 2 we introduce the four different astrophysical messengers. All of these have different advantages for detection, so multi-messenger astronomy combines different measurements for more conclusive information on the sources. Since this thesis is focussed on the UHECR-neutrino connection, a more detailed overview will be given for the observations of these two messengers as well as the prevailing interpretations of their main observed features. The relevant physical interactions and model assumptions required to interpret the flux and mass composition of UHECRs are explained in detail in Chapter 3. This includes the acceleration as well as interactions in the source, during extragalactic propagation and in the atmospheric air showers.

Our code for extragalactic propagation, PriNCe, will be presented in Chapter 4. It was developed from the ground up during the work leading up to this thesis. We focus on the main techniques used to accelerate the computation. The PriNCe code is employed in Chapter 5 to fit the UHECR energy spectrum and composition data in a comprehensive parameter scan. In this scan, we study the parameters of a generic source population, which is defined with the spectral index, the maximal rigidity, the cosmological source evolution and the mass composition as free parameters. We discuss the impact of different model choices for disintegration during extragalactic propagation as well as hadronic interactions during the air-shower development. From this parameter scan, we derive a prediction for the expected flux of cosmogenic neutrinos and its robustness under the model assumptions.

GRBs are discussed as possible sources of UHECRs and neutrinos in Chapter 6. We model them in a sophisticated multi-collision model that takes into account the radiation at different radii in the plasma jet of a GRB. In this model we discuss the impact of different assumptions on the GRB engine accelerating the plasma jet as well as the model for the hydrodynamic collisions occuring inside the jet. From this, we make predictions for the ejected neutrino and UHECR spectra and their dependence on the initial assumptions. Lastly, we discuss the expected neutrino emission from GRB17017A and its dependence on the observation angle in Chapter 7. This short GRB has received specific attention because its observation was the first time that an electromagnetic signal was observed in coincidence with a gravitational wave event. We discuss the impact of different assumption about the jet geometry on the expected neutrino flux.

Chapter 2

Multi-messenger astrophysics with UHECRs and neutrinos

At present, there are four messengers that can be detected in astronomy: electromagnetic radiation, cosmic rays, neutrinos and gravitational waves. Each of these messengers individually carries important information on their sources. However, some messengers are likely produced in common source environments. Especially the production of high-energy neutrinos is expected to be closely connected to the acceleration of UHECRs. Multi-messenger astronomy deals with exploiting these connections by trying to observe correlations between the different messengers. This thesis deals with models of multi-messenger production, specifically focussing on the UHECR-neutrino connection. This chapter gives an overview of the different messengers and how their combined information can lead to new discoveries in Section 2.1. This is followed by more detailed introductions of the current observations related to UHECRs in Section 2.2 and astrophysical neutrinos in Section 2.3. A short overview of the prevailing models for the origin of UHECRs and their neutrino constraints is given in Section 2.4.

2.1 Astrophysical messengers

Astronomy goes back several millenia and was based on observations with the bare eye for a long time. In more recent times, technical advances constantly pushed the range of the observable universe. New instruments allowed to observe an increasing range in photon-wavelengths. Additionally, more messengers became experimentally observable. A little more than 100 years ago, cosmic rays were first discovered and have since been detected at increasingly higher energies. More recently, neutrinos and gravitational waves completed the range of currently observable astrophysical messengers. Our knowledge of the Universe has not only been extended by the increasing number of observations, but also by combining their information, first in multi-wavelength and then in multi-messenger astronomy. This is because each of the different messengers has unique advantages:

Electromagnetic radiation: Still today, most of our information on the Universe comes from photon observations in different wavelengths. Most known sources were first identified by photon observations. The observable range in wavelength has been constantly extended by new technical developments and ranges from very low (radio) to very high (gamma rays) energies. Different wavelengths require a wide range of detection techniques: While radio detection requires ground based detectors with large area (Eq. ALMA [28], VLA [29]), X-rays and gamma-rays are best observed in space (e.g. Swift [30], Fermi [31, 32]), to avoid absorption by the atmosphere. Gamma rays in the TeV range can also be detected indirectly by ground based Cherenkov telescopes (e.g. H.E.S.S. [33], VERITAS [34] and MAGIC [35]). By combining the information from different wavelength bands, we get a more complete picture of astrophysical sources. As most detectors have a narrow field of view, this requires coordination of different experiments especially for observing transient sources. Observations are therefore often triggered by other wavelength bands or messengers that provide a wider field of view.

Cosmic rays: Cosmic rays are charged particles, mostly protons and ionised nuclei, that constantly hit our atmosphere. Since their discovery by Viktor Hess [36] in 1912, cosmic rays have been observed at increasingly high energies. With higher energies, these observations become increasingly complicated due to the steeply falling flux. At the highest energies, extensive detector arrays are needed to observe the air showers initiated by cosmic rays. Ultra-high-energy cosmic rays (UHECRs, $E_{\mathrm{UHECR}} \gtrsim 10^9$ GeV) are the highest energy particles ever detected, therefore carrying information on the most extreme astrophysical sources. As charged particles, cosmic rays are deflected by magnetic fields and lose directional information on their source over small distances. Diffusion also makes them lose temporal correlation with other messengers that travel on straight paths. The observations of UHECRs will be discussed in detail in Section 2.2.

Neutrinos: In most sources, the acceleration of high-energy cosmic rays will be accompanied by a target photon field. At high enough energy, cosmic rays will interact with these photons, producing secondary neutrinos and gamma rays. Neutrinos interact only through the weak force and can escape the source. Additionally, cosmic rays can produce neutrinos during extragalactic propagation by interacting with cosmic photon background. High-energy neutrinos are therefore a clear signature of hadronic sources. Astrophysical neutrinos are in principle the ideal messenger, since they are electrically neutral and barely interact and are therefore neither deflected nor absorbed during extragalactic propagation. However, the same properties are also a challenge to their detection. Most neutrino detectors work by observing the Cherenkov radiation from charged particles produced by the interaction of a neutrino with matter. At the highest energies, this requires a large natural volume of transparent medium. The largest astrophysical neutrino detector today is the IceCube Observatory, which is located in the Arctic ice [37]. It is the first detector to have observed a diffuse flux of astrophysical neutrinos [21]. The detection techniques and main results of the IceCube Observatory will be discussed in detail in Section 2.3.

Gravitational waves: Gravitational waves were first predicted theoretically by Albert Einstein shortly after he finished his theory of general relativity [38]. According to this theory, massive objects create a curvature in space-time around them. If these objects are accelerated, they induce changes in the space-time curvature, which can propagate as gravitational waves. Since the amplitudes of gravitational waves are normally very small, massive and rapidly accelerating objects are required to produce an observable signal. Colliding compact objects such as black holes and neutron stars are a primary source. Other sources are stellar core collapses and pulsars.

The Laser Interferometer Gravitational-wave Observatory (LIGO) detected the first gravitational wave signal in 2015 [39]. Gravitational waves carry unique information on the formation and evolution of compact objects, while other messengers carry information about acceleration, composition and interactions in the shocks formed by these objects. Furthermore, merging black holes produce no radiation other than gravitational waves. The sensitivity of gravitational wave detectors covers almost the entire sky. By combining the signals detected in detectors distributed around the world (currently LIGO Hanfort and Livingston [40], and Virgo [41]), the source locations can be determined with higher precision and trigger follow-up observations.

Each of the astrophysical messengers has unique advantages which, in combination, can reveal a broader perspective on the physics of their common sources. Revealing these connections requires effort from both the experiments, which need to coordinate observations as well as the theory, which need sophisticated models to predict the production of the different messengers.

A lot of effort is put into coordinating multi-messenger observations. This is especially important for transient sources. Many instruments that observe photons in different wavelengths have a small field of view. They therefore need external triggers to point them to targets of interest. Instruments suited to trigger such follow-up observations need a large field of view, optimally extending over the whole sky. This makes neutrino and gravitational wave observatories ideal monitoring devices. Conversely, directional information from other instruments can be used to look for signals in neutrinos and gravitational waves from specific directions, which would otherwise be below the detection threshold due to background. The gravitational-wave detectors LIGO [40, 44] and Virgo [41] have a public alert systems, which sends out real time triggers for follow up observations to over 80 other instruments. Similarly IceCube sends out alerts for neutrino multiplets and neutrino candidates with a high probability of astrophysical origin [45]. These trigger system have lead to two recent breakthroughs in multi-messenger astronomy.

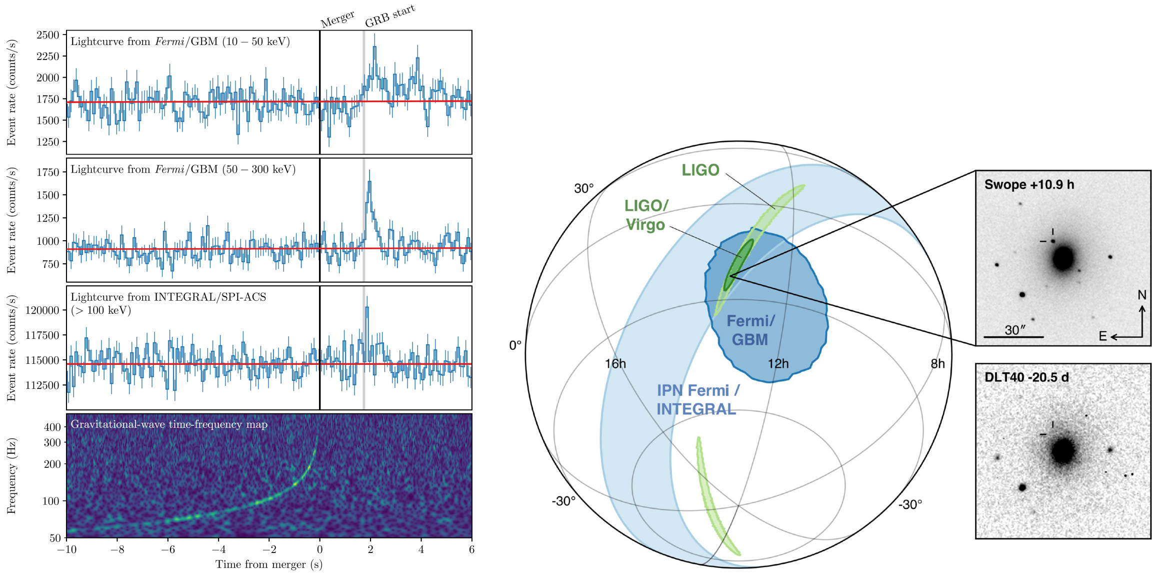

Figure 2.1: Multi-messenger observation of GW170817 and GRB17017A. **Left:** The combined frequency map of the gravitational wave signal observed the two LIGO detectors is shown in the bottom panel. The corresponding GRB light curves observed in Fermi-LAT, Fermi-GBM and INTEGRAL are shown in the top three panels. There is a delay of 1.7 s between the gravitational wave signal and the GRB. Reprinted from Abbott et al. [42] under CC BY 3.0 license. **Right:** Localization of the gravitational wave (green) and gamma-ray (blue) signals. The precise localization by the combined data from LIGO and Virgo allowed for follow up observations that identified the host galaxy NGC 4993. Its location observed in the optical range is shown as an inset. Reprinted from Abbott et al. [43] under CC BY 3.0 license.

In August 2017 the LIGO collaboration send out an alert for the gravitational wave event GW170817 [46, 47], 1.7s later the short gamma-ray burst GRB170817A was detected by the Fermi Gamma-ray Space Telescope and the INTErnational Gamma-ray Astrophysics Laboratory (INTEGRAL) [42]. This triggered a wide range of follow-up observations in different wavelengths, which eventually identified the host galaxy NGC4493 [43]. Fig. 2.1 shows the observed gravitational wave signal and GRB light curves, and their localization. Specifically the precise localization by the gravitational wave compared to the GRB signal allowed for the swift and extensive follow-up campaign. We will discuss the neutrino production for this event in detail in Chapter 7. In September 2017, the IceCube collaboration detected a 290 TeV neutrino event called IceCube-170922 [48], which caused triggers to be sent out. Shortly after, the blazar flare TXS0506+056 was detected by the Fermi-LAT and MAGIC collaborations in a consistent direction [49]. Using the directional information on the blazar location to constrain backgrounds, an excess of neutrinos was detected between September 2014 and March 2015 at a significance of 3.5 σ [50]. However, there was no notable gamma-ray activity in the same period, making it hard to reconcile this neutrino excess with theoretical models [51, 52].

These two events were, however, not the first multi-messenger observations. In 1987, a burst of neutrinos from the supernova SN1987A was detected in the Kamiokande-II [53], IBM [54] and Baksan [55] neutrino detectors. This was the first observation of neutrinos from a source outside of our solar system. This lead to the development of the Supernova Early Warning System [56] consisting of different neutrino detectors, which will send out automatic alerts in the case of a close-by supernova. Other networks exist for specific source types like the Gamma-ray Burst Coordinates Network (GCN), that developed from the BATSE coordinates distribution network (BACODINE) [57]. The Astrophysical Multimessenger Observatory Network (AMON) [58, 59] is further expanding this idea to include all cosmic messengers. By combining and analyzing the sub-threshold data from multiple messengers, it can identify source candidates that would not be identifiable from single messengers.

Observational multi-messenger astronomy aims at identifying point sources or source classes by finding directional correlation between different messengers. Direct multi-messenger observations are unlikely to include UHECRs, which lose spatial and temporal correlation with their sources quickly by magnetic diffusion. However, even without these direct correlations, information can be gained by considering the diffuse fluxes of different messengers in concurrence. Remarkably, the energy budget of the diffuse fluxes of UHECRs, cosmic neutrinos and unresolved gamma-rays is on a comparable level. This suggests that they might be produced by a single dominant source class. Theoretical models based on the UHECR-neutrino connection are significantly more constrained than ones considering only a single messenger. While a model of the UHECR flux must not necessarily explain the neutrino flux as well, it should at least not exceed the observational limits. In this thesis, we focus on the UHECR-neutrino in models of the diffuse flux. Their observations will therefore be described in more detail in the following two sections. Their theoretical connections will be discussed in more detail in Section 2.4.

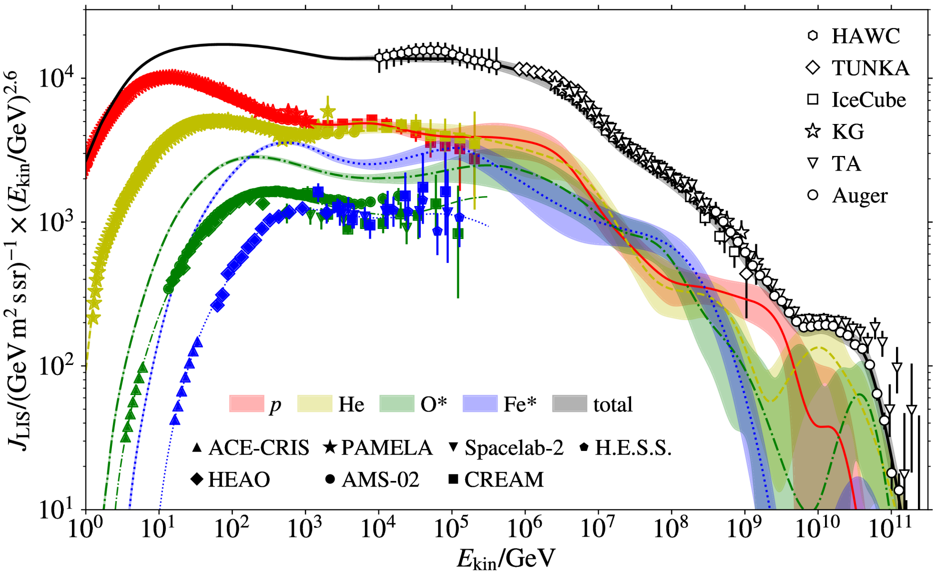

Figure 2.2: All-particle flux of cosmic rays (black) over the full observed energy range. The contribution of different (element groups) are shown for protons (red), helium (yellow), the oxygen group (green) and the iron group (blue). The curves and uncertainty bands are from a Global Spline fit to the data as described in Dembinski et al. [60]. Data for composition at high energies is not shown for clarity. Reprinted from Dembinski et al. [61] under CC BY 4.0 license.

2.2 Ultra-high-energy cosmic rays

The spectrum of cosmic rays stretches over several orders of magnitude from 1 GeV to the end of the observed spectrum at $\sim 10^{11}$ GeV, the highest energy particles ever observed. The flux is a steeply falling power-law with little distinct features. This constitutes a challenge for detection. At lower energies, cosmic rays can be directly observed by either balloon or satellite based experiments. These include ACE-CRIS [62, 63], HEAO [64], PAMELA [65, 66], AMS-02 [67–69], Spacelab-2 [70, 71] and CREAM [72–74]. Above $\sim 10^5$ GeV however, the flux of cosmic rays decreases to less than one particle per square meter and year making direct observations unfeasible. Instead, the required exposure is reached by observing extensive air showers initiated by cosmic rays in the atmosphere. Examples are ARGO-YBJ [75], TUNKA [76], IceTop [77] and KASCADE-Grande (KG) [78], AGASA [79] and HiRes [80]. Cherenkov telescopes like H.E.S.S. are also sensitive to cosmic ray air showers [81]. The two largest cosmic ray detectors operating today are the Telescope Array (TA) [9–11] and Pierre Auger Observatory (Auger) [6–8] observing up to the highest energies ever detected.

The spectrum of cosmic rays over the full energy range is shown in Fig. 2.2, with the data compiled from the measurements of different experiments. It is given as the number of particles per energy, unit area, time and solid angle. To emphasize the spectral features, the spectrum is multiplied by $E^{2.6}$. To model the flux and composition over the whole energy range a global spline fit was done by Dembinski et al. [60] for four different mass groups, taking into account the systematic uncertainties of the experiments. The best fit and uncertainties for each mass group are shown as colored bands.

There are a few distinct features that have been given names after their similarity to a human leg: At lower energies the flux follows a power-law with a spectral index of $E^{-2.7}$ up to a steepening at $\sim 4 \cdot 10^6$ GeV, referred to as the *knee*. A less pronounced feature occurs as a *second knee* at $\sim 10^8$ GeV. Finally, a hardening at $\sim 5 \cdot 10^9$ GeV is called the *ankle*. The observed spectrum ends in a sharp cutoff around $5 \cdot 10^{10}$ GeV. The uncertainties of composition measurements increase strongly with higher energy. At lower energies, direct detection methods can distinguish particles by their charge Z and therefore measure the flux of CRs in terms of individual elements. Indirect detection methods above 10^4 GeV can instead only distinguish roughly between different mass groups. Cosmic rays at energies $\sim 10^9$GeV are typically referred to as UHECRs. We will focus in their measurements in the following.

2.2.1 Extensive-air-shower detection

The two experiments measuring the spectrum of UHECRs today are the Pierre Auger Observatory (Auger) [7, 8] and the Telescope Array (TA) [11]. The main detectors in both experiments comprise a surface detector array (SD) and a fluorescence detector (FD). The basic geometry of the extensive air shower detection in Auger is shown in Fig. 2.3 (left panel). These two detection principles complement each other: the SD is an array of individual detectors for charged particles, either scintillation detectors or Cherenkov tanks. It measures the time structure and lateral distribution of the electrons and muons in the air-showers arriving at ground level. This leads to a large exposure due to a duty cycle of almost 100%. However, the SD can only observe a cross section of the shower at ground level. The data analysis therefore relies on models of the air-shower development to reconstruct properties of the initial particle such as its energy. The FD operates by observing the longitudinal profile of the fluorescence light produced by interaction of the shower particles with the atmosphere. This light is produced by the de-excitation of nitrogen atoms that are exited by high-energy photons from the air-shower. The FD can only operate on clear, moonless nights, severely limiting its duty cycle to $\sim$ 10%. Its energy measurement is almost calorimetric, as the fluorescence light is proportional to the energy deposited in the atmosphere by the air shower. This method therefore gives a more accurate estimate of the initial UHECR energy than that from SD detections.

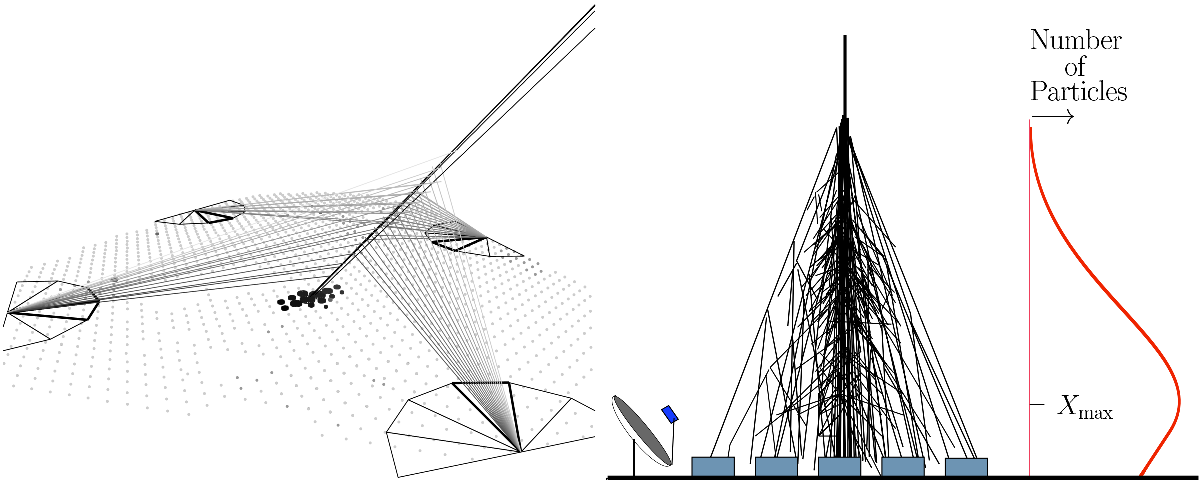

Figure 2.3: Depiction of the hybrid air shower detection method. **Left:** Three dimensional depiction of an event reconstructed in the Pierre Auger Observatory. The SD array (grey dots) detects a cross section of the shower at ground level (bold black dots). The atmosphere above the SD is overlooked by four FD stations that observed the longitudinal development of the air shower. Reprinted from [6] under CC BY-NC-ND 4.0 license. **Right:** Schematic illustration of the air shower development. The initial UHECR initiates a cascade driven by hadronic and electromagnetic interactions. The atmospheric depth where the particle number reaches its maximum is labelled X_{max} Reprinted from [15] with permission from Elsevir.

By exploiting the subset of 'hybrid' events, which are observed in both the SD and FD, the energy scale of the SD can be calibrated to that of the FD. The hybrid detectors therefore combine the advantages of both techniques. A large data-sample is obtained due to the exposure of the SD with the model independent energy scale of the FD. Additionally, it provides a handle on the composition of the primary UHECR. It is usually inferred from the lateral profile of the shower development, more specifically the atmospheric depth where the shower reaches its maximum X_{max}, see Fig. 2.3 (right panel) for an illustration. This lateral shower profile can only be observed in the FD. Still, the mass composition can also be inferred from SD measurements through sensitive variables such as the difference in rise time between individual SD-stations [82]. These, however, have much higher systematic uncertainties than the FD measurements.

TA is located in Utah and has been operating since March 2008. The SD consists of 507 scintillation detectors, which are arranged over an area of 700 km^2. 48 fluorescence detectors distributed over three stations compose the FD. TA's predecessor, the High Resolution Fly's eye (HiRes) [83, 84] was at the same location, but relied only on the FD technique. Auger, located in west Argentina, has been taking data since 2004 and was fully completed in 2008. The SD is composed of 1660 Cherenkov water tanks, each consisting of 12 tons of clear water

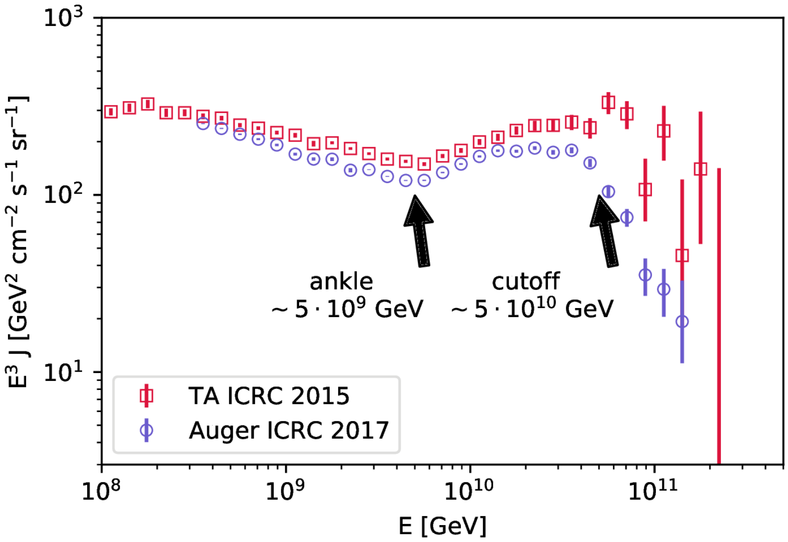

Figure 2.4: Spectrum of UHECRs as measured by the Pierre Auger Observatory [88] and the Telescope Array [89] respectively. There is a systematic shift of the energy-scale between the two spectra of about 10% [90]. When correcting for this, the two spectra are in good agreement, with the exception of an access of events in the TA spectrum at the cutoff.

monitored by three photomultiplier tubes. They cover an area of about 3000 km^2 making Auger the largest cosmic-ray detector on Earth. The FD consists of four stations each equipped with six fluorescence detectors.

Both experiments currently have major upgrades planned. For AugerPrime [85] each of the water Cherenkov tanks will be refurbished and equipped with an additional scintillation detector. This will improve the separation between electrons and muons to improve mass composition reconstruction. Additionally a radio antenna array is planned, which will further improve the shower reconstruction [86]. For the TA×4 upgrade [87], the surface array will be expanded and and new fluorescence detectors installed. This will increase the area to roughly the same as that of Auger, allowing for measurements with higher statistics in the Northern Hemisphere.

2.2.2 UHECR spectrum and composition

The spectrum of UHECRs has been measured by both TA and Auger using the hybrid approach. The ‘combined’ data-sets, i.e. combining SD and FD measurements, of either experiment are shown in Fig. 2.4. They both show the two features described in Section 2.2: The *ankle*, a hardening at $\sim 5{\cdot}10^9$ GeV, and a cutoff around an energy of $\sim 5{\cdot}10^{10}$ GeV. There is a systematic shift between the spectra of the two experiments. Using the *ankle*, in which both experiments have high statistics, for calibration, the relative shift in the energy scale can be determined to $\sim 10\%$ [90]. This shift is fully consistent with the systematic uncertainties reported on the energy scale by Auger (13%) and TA (20%). Accounting for this shift in energy the *ankle*-region is in good agreement. However, there seems to be an access of the TA data over the Auger spectrum in the cutoff region.

A recent analysis suggests that this is a physical feature rather than a fluctuation. The measurements are in good agreement when comparing only those events observed in declination bands common between the two experiments [90]. Anisotropy analyses provide further indications for this. The TA collaboration has been reporting a hotspot in the northern hemisphere at a declination unobservable by Auger [91]. This hotspot might be the source of the excess at high energies.

To infer the composition of the primary UHECRs, the most sensitive observable is the atmospheric depth of the shower maximum $X_{\max}$ (measured in g cm^{-2}). This is the depth in the atmosphere, where the shower deposits most of its energy, as illustrated in Fig. 2.3. Heavy elements develop showers earlier in the atmosphere than proton showers, leading to smaller $X_{\max}$, while also having smaller shower-to-shower fluctuations in $X_{\max}$. This can be understood qualitatively by viewing the nuclei as a superposition of nucleons ignoring binding energy [92]. In this picture, the number of particles after the initial interaction is proportional to the mass number A, leading to a scaling of $X_{\max}$ with $\ln(A)$. For a quantitative analysis, hadronic interaction models are needed, which extrapolate from the LHC scale to the energy scale of the initial interaction. More details on these models will be given in Section 3.3.

As the shower-to-shower fluctuations in $X_{\max}$ are larger than the resolution of the detectors, the composition of an UHECR cannot be inferred for a single shower. Instead, the average composition is inferred from the distribution of $X_{\max}$ for showers of similar energies. This is typically presented in terms of the first two moments $\langle X_{\max} \rangle$ and $\sigma(X_{\max})$. While $\langle X_{\max} \rangle$ is sensitive to the mean mass, $\sigma(X_{\max})$ depends both on the mean mass and the variance in mass. The Auger collaboration provides the highest statistics on the composition of UHECRs. Its measurements are presented in Fig. 2.5 and compared to calculations from different hadronic models for a range of elements. The average composition is light around the *ankle*, with proton or helium like composition, and becomes increasingly heavier with energy. Even though the inferred mass is model-dependent, the trend is the same for all hadronic models. The data on $\sigma(X_{\max})$ shows a similar trend, which also indicates relatively pure composition at the different energies. The Auger collaboration has also analysed the composition using the SD [82]. Due to the higher duty-cycle this analysis can be extended to higher energies than the FD. These measurements indicate that the composition is getting lighter again above 40 EeV.

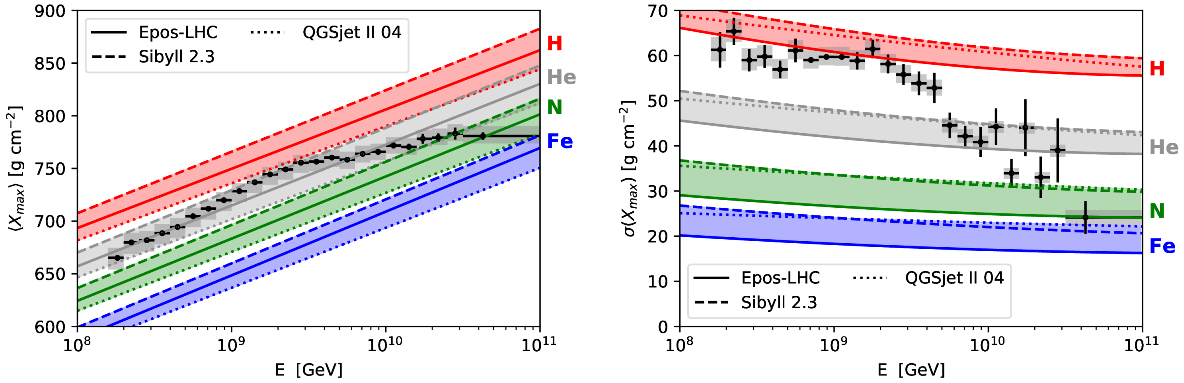

Figure 2.5: First two moments of $X_{\max}$ for UHECRs as reported by the Pierre Auger Observatory [93]. To infer the mass of the primary, the data is compared to different hadronic interaction models. The $\langle X_{\max} \rangle$ indicates a composition that is helium-like at lower energies and becomes lighter up until the *ankle* at $\sim 3 \cdot 10^9$ GeV. At higher energies the composition becomes increasingly heavier again. The same trend is visible in the $\sigma(X_{\max})$ data. This data set additionally implies that there is relatively little superposition of masses above the *ankle*. Plot adapted from Heinze et al. [3]

The TA collaboration has published a similar analysis of their observed $X_{\max}$ distribution, which indicates light, helium-like composition up to the highest energies [94]. At first glance, this seems to be in conflict with the Auger result. However, TA has much lower statistics and relies on the older hadronic interaction models QGSjetII-03 and QGSjetII-04. A direct comparison is complicated because TA does not unfold their $X_{\max}$ measurements in terms of detector response. By running the composition measured by Auger through the TA analysis tools, the two collaborations have concluded that the measurements are actually in good statistical agreement [95]. This means that TA simply cannot distinguish an Auger-like composition from a light composition due to its lower statistics.

Additional information comes from the arrival direction of UHECRs. Directional correlations to sources are difficult to measure due to magnetic deflections. However, anisotropies can be observed on larger scales. An anisotropy analysis performed by the Auger collaboration finds a dipole anisotropy at the highest energies [12]. Auger also analysed correlations with source catalogues, allowing for magnetic deflections. The most significant correlation is found for Starburst Galaxies with $\sim 4\sigma$ [14]. In this thesis, we do not take into account directional data and instead analyse only the diffuse spectrum and composition. We will focus on interpreting the spectrum and composition measurements from Auger, as they provide the higher statistics.

2.3 Astrophysical neutrinos

The identification of the sources of UHECRs is significantly complicated by them being deflected in magnetic fields and interacting with photon backgrounds. This limitations can in principle be overcome by the observation of astrophysical neutrinos. These are dominantly produced as secondaries from photo-hadronic interactions of charged UHECRs on photon-target fields. At threshold, these interaction are described by the Δ^+-resonance

$$p + \gamma \to \Delta^+ \to \begin{cases} n + \pi^+ & \text{with } 1/3 \text{ BR} \\ p + \pi^0 & \text{with } 2/3 \text{ BR} \end{cases} . \tag{2.1}$$

Neutrinos are produced in the decay chain of π^+, while the π^0 decays mostly into photons. These details of these photo-hadronic interactions will be discussed in Chapter 3. As neutrinos interact only through weak interactions, they escape even dense sources and can travel over cosmological distances without interacting. Additionally, they are electrically neutral and therefore travel on a straight trajectory to the observer. Observed neutrinos therefore point directly to their production site. While UHECRs and gamma rays are limited to the local Universe due to interactions with extragalactic photon backgrounds, neutrinos can in principle probe as far as the early Universe.

While these properties make neutrinos an ideal extragalactic messenger, they also make them very difficult to detect. The main technique to detect neutrinos is observing the Cherenkov light of a charged particle, which is produced by the weak interaction of a neutrino with matter. This requires a very large transparent medium. For lower energy detectors such as the Super-Kamiokande detector [96], enough ultra-clear water can be produced artificially. This allows for the detector to be placed underground to reduce backgrounds. Higher energy detectors need to exploit a natural occurrence of transparent medium. The ANTARES detector uses the water of the Mediterranean Sea [97]. It will be succeeded by the KM3NET, which will consist of the two sub detectors ARCA and ORCA [98]. Currently, the biggest neutrino detector is IceCube[37], which instruments a volume of $\sim 1\ \mathrm{km}^3$ of transparent Arctic ice. An upgrade called IceCube-Gen2 is currently being planned [99] which will increase the instrumented volume.

Secondaries from neutrinos in principle also produce air-showers, much like cosmic rays. At large zenith angles and for Earth-skimming events, the background can be reduced enough to make air-shower experiments like Auger sensitive to neutrinos around 1 EeV [100]. However, no neutrinos have been detected by Auger yet. To push the sensitivity to higher energies, future experiments need to instrument even larger volumes. One promising approach is to use radio antennas, which are cheap to produce. Both ARA [101] and ARIANNA [102] aim at observing

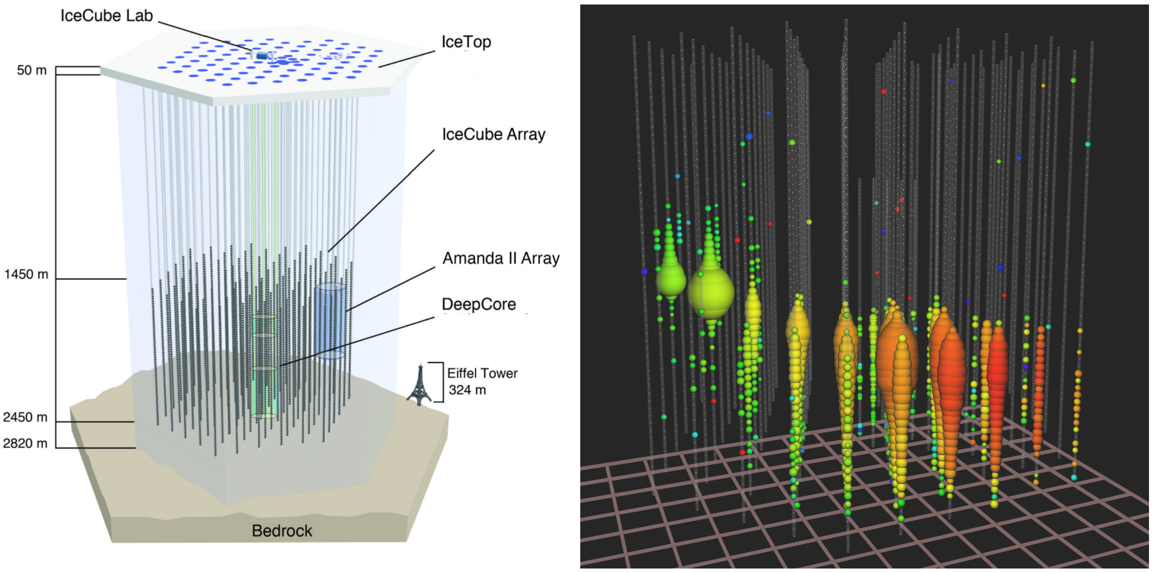

Figure 2.6: **Left:** Schematic depiction of the IceCube Observatory. The detector consists of 86 strings drilled into the ice that hold the individual digital optical modules (DOM). A detector array at the surface (IceTop) detects cosmic rays and can be used to veto events. **Right:** Cherenkov signal produced in the detector by an upgoing muon track with a reconstructed energy of 2.6 PeV. The energy deposited in each DOM is represented by the size of the colored dots. The color indicates arrival time from red (early) to purple (late). Reprinted from Ahlers and Halzen [105] with permission from Elsevier.

radio signals from neutrino secondaries in the Arctic ice. GRAND [103] aims at using radio antennas to monitor a large volume of the atmosphere, making it mainly sensitive to very-high-energy Earth-skimming tau neutrinos. POEMMA aims at observing the same type of event but using a satellite looking down on the atmosphere [104]. Since they observe extensive air showers, POEMMA and GRAND will also be sensitive to UHECRs.

2.3.1 The IceCube neutrino observatory

IceCube is a neutrino observatory located at the geographic south pole [45]. It consists of 5000 digital optical modules (DOM) deployed along 86 strings over a km^3 of Arctic ice. A schematic illustration of IceCube is shown in Fig. 2.6. A neutrino is detected when it interacts inside the ice and produces a charged lepton. The lepton will carry most of the energy of the primary neutrino and produce Cherenkov radiation inside the transparent ice. An overview of the reconstruction methods in IceCube is given in Aartsen et al. [106]. A recent review of the results and their interpretation can be found in Ahlers and Halzen [105].

The signals observed by IceCube can be categorized into two main topologies, *showers* and *tracks*. A reconstructed muon track is shown Fig. 2.6 (right panel). A track occurs when the neutrino produces a muon. Due to its mass, the muon can travel several kilometers inside the ice, depositing some of its energy along the way. Tracks provide a good angular resolution down to 1 degree. However, they deposit only a fraction of their energy inside the detector, making energy reconstruction challenging. Showers are produced by electron or tau leptons. Electrons lose their energy quickly in interactions due to their small mass, while tau leptons decay almost immediately. Both cases produce a spherically shaped electromagnetic cascade inside the ice, which is detected by the DOMs. If the whole event is contained inside the detector, its energy can be reconstructed with a resolution of up to 15%. However, the spherical shape provides a weak angular reconstruction of at most 10 degrees. A special topology is the so-called *double-cascade*, two cascades connected by a track. This signal is expected if a tau travels enough distance before decaying to be resolved by the detector. Neutrino oscillations necessitate some abundance of tau neutrinos and therefore double cascade events are expected. In fact, two candidates for double cascade events have recently been identified [107].

2.3.2 The diffuse flux of astrophysical neutrinos

The large expected background from atmospheric muons provides a strong challenge to the detection of neutrinos. In order to establish the detection of an astrophysical flux, this background has to be reduced significantly. It is overcome by different techniques. A surface detector, called IceTop, measures the flux of cosmic rays at the surface of the ice [106]. Its primary intent, however, is to veto atmospheric muons for the main detector. Additionally, cuts can limit the background significantly by restricting to events that start inside the detector (*high energy starting events*, HESE) or muon tracks from the Northern Hemisphere (*through going muons*, TGM). The fluxes from the two samples are shown in Fig. 2.7 (left panel).

The first evidence for cosmic neutrinos was established by the HESE analysis [21, 22]. It separates the detector into an outer veto layer and an inner signal region. By considering only events that start inside the inner detector, the expected atmospheric background is reduced to a few events per year. For the 6 year sample 82 HESE events have been observed. An atmospheric origin of the observed flux can be rejected at 8σ. A single-power-law fit to the events yields a spectral index of $2.92^{+0.29}_{-0.33}$ [109]. The TGM analysis reduces the atmospheric background by restricting the direction to those in which atmospheric neutrinos are mostly absorbed in the Earth. At energies above 100 TeV, the observed flux of TGM events shows an excess over the expected atmospheric background. Even with large uncertainties in the extrapolation of the

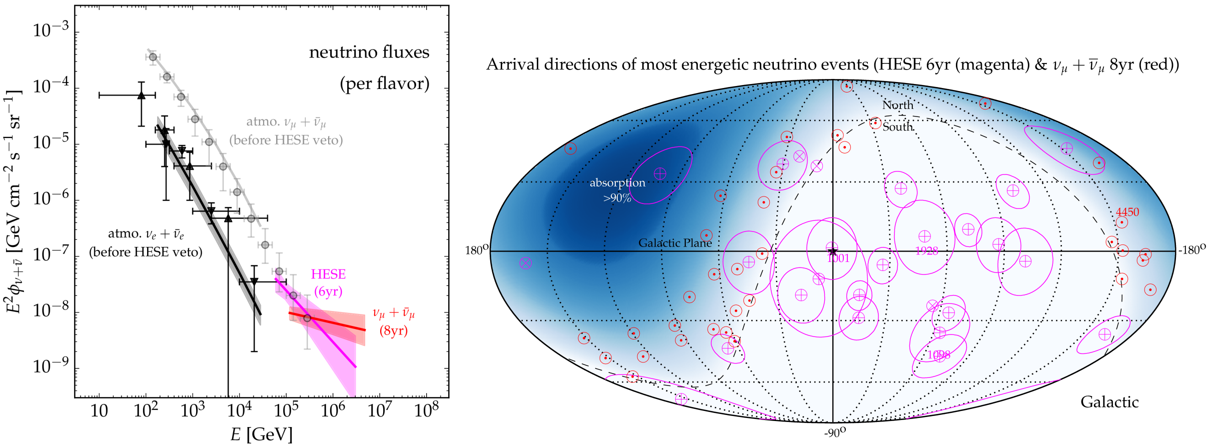

Figure 2.7: Diffuse flux and arrival directions of astrophysical neutrinos. **Left:** The flux of astrophysical neutrinos as measured by IceCube in different samples: High energy starting events (HESE, magenta) [108] and through-going muons (TGM, red) [109]. Both fluxes are shown as their power-law best fit with 1σ uncertainties. The steep atmospheric backgrounds are shown in grey and black. **Right:** Arrival direction of events with energies above 200 TeV (TGM, $\odot$) and 100 TeV (HESE, $\otimes$ and $\oplus$). For HESE-cascades the median angular resolution is shown as thin circles. The dashed line indicates the horizon and the star ($\star$) the Galactic center. The blue shading indicates the area where absorption inside the Earth becomes important for 100 TeV neutrinos. The energies of the four highest-energy events of each sample are shown as numbers. Reprinted from Ahlers and Halzen [105] with permission from Elsevier.

atmospheric background, an astrophysical flux was established with a significance of 6.7σ. In 8 years, 36 such TGM events have been observed at energies above 100 TeV. The corresponding flux is consistent with a spectral index of 2.19 ± 0.10 [108, 110].

Though the spectral index of the HESE flux is softer than the flux of the TGM analysis, they are still consistent above 200 TeV within 2σ [109]. However, the HESE flux at energies below 100 TeV shows an excess over the TGM flux extrapolated to these energies. This excess is difficult to explain by atmospheric background due to the strong background veto of the HESE sample. Considering instead a two-component power-law fit for the HESE flux yields one component with $\gamma = 3.7$ at low energies and one above 200 TeV with $\gamma = 2.3$, where the high energy component is consistent with the TGM results [109]. This might point towards multiple components in the astrophysical neutrino flux [111].

The arrival direction of the eight-year TGM and six-year HESE samples are shown in Fig. 2.7. The angular resolution of each event is indicated by a circle. TGM events appear clustered around the horizon. At lower angles, they are increasingly absorbed while propagating through the Earth, while at higher angles the background veto is not efficient enough. The HESE sample is similarly affected by absorption in the Earth. When accounting for these effects, the distribution of events is consistent with an isotropic distribution, which indicates an extragalactic origin. However, no correlation to a specific source class has reached a significant level yet [109, 112]. This absence of correlations leads to stacking limits on the neutrino flux from different source types, e.g. gamma-ray bursts (GRB) [25, 113] and active galactic nuclei (AGN) [23]. The shape of the observed flux is also inconsistent with cosmogenic neutrinos, which would peak at even higher energies [26].

2.4 Origin of UHECRs and neutrino constraints

In this section, we give a brief overview of the main astrophysical models that explain the UHECR-spectrum and their constraints imposed by the observed astrophysical neutrinos. For more detailed reviews see e.g. Refs. [15, 105, 114].

The fact that the spectrum of cosmic rays is relatively featureless over many orders of magnitude implies that it is dominated by a small number of source classes. The prominent features are then attributed to the transition between different components. The power-law implies that the acceleration mechanism at the sources is likely non-thermal. Cosmic-ray hadrons will interact with the photon fields when accelerated at the sources, producing secondary pions, which decay producing neutrinos and gamma rays. At threshold, neutral π^0 and charged pions π^+ are produced through the Δ^+-resonance, Eq. (2.1), at a fixed radio of $\pi^0/\pi^+ = 2$, the energy going into either particle species is therefore roughly the same. In fact, the energy budgets of the UHECR-flux, cosmic neutrinos and gamma rays are on a similar level. It is therefore natural to assume that these particle come from the same sources. Multi-messenger models use this connection to constrain possible source models.

The most common models of cosmic rays up to the *knee* ($\sim 3 \cdot 10^6$ GeV) explain the observed flux by acceleration at supernova remnants (SNR) in our Galaxy [115]. The softening of the spectrum is then due to the sources reaching their maximal acceleration energy, which is determined by their magnetic confinement in the source. Alternatively, the *knee* can be explained by the propagation of cosmic rays in the galaxy [116]. At PeV scale energies, confinement in galactic magnetic fields becomes inefficient. The leakage of cosmic rays from the galaxy can then explain the steepening. This model, however, requires galactic sources that are able to accelerate particles to energies well above the *knee* energies. In both cases the *knee* energy will scale with the particle charge. Nuclei would be more efficiently confined due to their larger charge Z, with the maximal energy scaling proportional to Z. Iron cosmic rays then reach energies 26 times higher, which coincides with the *second knee* ($\sim 10^8$ GeV). This also supported by the composition measurements which show a transition from light to heavy between the *first* and *second knee*.

The origin of cosmic rays above the *second knee* is less clear. It is generally assumed that cosmic rays above the *ankle* ($\gtrsim 5{\cdot}10^9$ GeV) have an extragalactic origin. There are no known sources in our Galaxy that would be capable of reaching such high energies. Secondly, it is difficult to reconcile the observed isotropy of UHECRs with a galactic origin. The flux is more naturally explained by luminous and abundant extragalactic sources. AGN [117–120] and GRBs [121–124] have been considered as such sources. Starburst galaxies (SBG) [125–127] are also luminous enough to sustain the flux of UHECRs.

A smoking gun would be a correlation with the astrophysical neutrinos observed by IceCube, which has, however, not yet been observed. IceCube therefore provides stacking limits for the neutrino flux from luminous sources such as AGN and GRB, which constrains their contribution to the UHECR flux [25, 110]. Due to the high radiation densities in high-luminosity sources, they are expected to produce many neutrinos. Additionally, the interactions with photon fields limit the acceleration of UHECRs. Therefore, less luminous and more abundant source classes have gotten attention in the more recent literature. Since these source are partly below the detection limit in gamma rays, they are less constrained by neutrino data. Such candidates are tidal disruption events (TDE) [128–131] and low luminosity GRBs (LL-GRB) [132–134]. These sources are most abundant in the local Universe. Therefore they should show stronger clustering in neutrinos, meaning they can in principle be constrained by neutrino multiplet observations [135].

While both a galactic as well as an extragalactic component in the cosmic-ray spectrum are established, the exact transition is still unclear. Before the Auger composition measurements, the observation of *ankle* and *cutoff* were seen as evidence for UHECR protons. A cutoff at $6 \cdot 10^{10}$ GeV was predicted theoretically from interactions of protons on the Cosmic Microwave Background (CMB) by Greisen [136], Zatsepin and Kuzmin [137], about 50 years before its first detection [80]. This effect is referred to as the *GZK-cutoff*. Additionally, energy losses from photo-pair-production peak for protons at energies coinciding with the *ankle*. This pure proton model is referred to as the *dip-model* [138–141]. It is now strongly disfavoured by the Auger X_max measurements, which indicate an increasingly heavy composition above the ankle. Reconciling the observed X_max with a proton dominated composition would require significant modifications to the current hadronic interaction models [142]. Even then, the dip-model is still strongly challenged by the non-observation of cosmogenic neutrinos and gamma rays [5, 143, 144].

For UHECR nuclei, the spectrum and composition between the *ankle* and *cutoff* can be explained by a transition from light to heavy nuclei similar to the transition between the two *knees*. The general properties of such a component are often studied in models assuming identical and isotropic sources, which follow a power-law spectrum with a rigidity dependent cutoff [123, 145–149]. As these 'generic' sources are not associated with a specific cosmological source evolution, it is typically treated as an additional free parameter. A typical choice is $(1+z)^m$, where m can be both positive and negative. This choice covers most known source evolutions up to $z = 1$, which is the maximal redshift from which UHECRs reach us due to energy losses. UHECR nuclei are much less efficient at producing cosmogenic neutrinos than protons, as the production efficiency depends mainly on the energy per nucleon [150–155]. Most UHECR nuclei models are therefore unconstrained by the limits on cosmogenic neutrinos. Also, cosmogenic neutrinos do not interact during extragalactic propagation, and therefore their predicted flux is affected by the choice of source evolution at higher redshift [5, 27, 150–153, 156–159]. A subdominant proton component that would not significantly affect the Auger measurements, could still significantly contribute to the cosmogenic neutrinos flux [154, 155, 159]. One of the goals in the development of future radio detectors for neutrinos is to aim to detect cosmogenic neutrinos from nuclei, which requires large sensitivities at high energies [102, 103, 159].

Fits to the spectrum and composition in propagation models show a preference for hard spectra and local sources [147, 148, 160], as the Auger $\sigma(X_{\max})$ shows little superposition of masses. When considering magnetic fields, the spectra can be softer, as lower energy cosmic rays are trapped in these fields and experience a *magnetic horizon* [149, 161]. In these models, the *ankle* has to be explained by the transition to an additional, galactic component with heavier composition. In more realistic source-propagation models, the transition across the ankle can be partly explained by disintegration inside the sources [123, 134, 159, 162–165]. If the source model can effectively trap lower-energy UHECRs magnetically, it acts as a high-pass filter, producing the required hard spectra. Disintegration of UHECR on the photon fields then leads to a second component of light cosmic rays with a softer spectral index, which can partially account for the *ankle*. Still, an additional galactic component is usually needed to get a good fit to the observed spectrum and composition below the *ankle*.

Chapter 3

Acceleration, interactions and propagation of UHECRs

UHECRs are expected to be accelerated in astrophysical sources from where they have to propagate through extragalactic space before they can be detected by the air showers they initiate in the Earth's atmosphere. The sources of UHECRs and the exact acceleration mechanisms have not yet been uniquely identified. Depending on the assumed source class, different processes can lead to the observed non-thermal spectrum. Additionally, the acceleration mechanism competes with energy losses due to interactions with target photon fields or matter, that also produce secondary particles such as neutrinos and gamma rays. During extragalactic propagation, UHECRs interact with cosmic photon backgrounds and are deflected in magnetic fields. Lastly, the air shower development is driven by hadronic interactions with air atoms. In all of these environments, there are model uncertainties due to incomplete knowledge of the interaction cross sections and target particle fields.

The main observations related to UHECRs and astrophysical neutrinos are described in Chapter 2. This chapter gives an overview of the model assumptions that enter in each of the three parts. There is a wide range of possible UHECR source candidates with different source physics that cannot be covered in this introduction. We focus on giving an overview of the basic acceleration mechanisms and general constraints in Section 3.1. The overview of extragalactic propagation in Section 3.2 is more comprehensive. We introduce all the relevant physical processes that are needed for the implementation in the PriNCe code, which is described in Chapter 4. For the air-shower development, we give an overview in Section 3.3 of how the different hadronic models change the inferred mass composition.

3.1 Acceleration and possible source candidates

The exact mechanism through which UHECRs are accelerated to their extremely high energies is still unknown. The observed power-law spectrum indicates that the acceleration happens through some form of magneto-hydrodynamic process. The exact details depend strongly on the physical environment at the source. There are two main types of processes currently considered:

stochastic and direct acceleration. For a detailed review of these processes and possible source classes see e.g. Anchordoqui [15]. This section summarises only the main ideas.

The basic principle behind stochastic acceleration is that particles gain energy gradually by scattering off magnetised plasmas or gas clouds. This idea was initially developed by Enrico Fermi [166, 167] and is therefore usually also referred to as 'Fermi acceleration'. In his initial publication, Fermi considered scattering of charged particles in between magnetised gas clouds, leading to an average energy gain of $\langle \Delta E/E \rangle \propto \beta_{\mathrm{cl}}^2$ (2nd order Fermi acceleration), where β_{cl} is the average cloud speed. This process is however inefficient since $\beta_{\mathrm{cl}}^2 \ll 1$. Therefore, diffusive shock acceleration (or 1st order Fermi acceleration) is more commonly considered for the acceleration of UHECRs. It takes place at collisionless plasma shocks. These shocks are formed by plasma discontinuities that travel at supersonic speeds. The shock separates the plasma into two parts: the upstream region ahead of the shock and the shocked downstream region. Particles on either side of the shock will diffuse in magnetic turbulences until they are either deflected across the shock front or escape from the shock region. The plasma flow is converging towards the shock in the respective rest frame on either side. Therefore, particles that are deflected across the shock front always gain energy in head-on collisions resulting in a first order energy $\langle \Delta E/E \rangle \propto \beta_{\mathrm{sh}}$. An unsolved problem is that diffusive shock acceleration requires a pre-accelerated component. Only particles with non-thermal energies decouple from the plasma and can cross the shock front to enter the acceleration-cycle [168].

Both types of stochastic acceleration naturally produce power-law spectra as the energy gain per step $\xi = \langle \Delta E/E \rangle$ competes with a chance to escape the acceleration region P_{esc} [169, chap 3.2]. After n scattering steps, a particle will have an energy of $E_n = E_0(1+\xi)^n$, conversely the number of encounters to reach an energy of E_n is $n = \ln(E_n/E_0)/\ln(1+\xi)$. The chance that a particle encounters n collision before escaping is $P_n = (1 - P_{\mathrm{esc}})^n$. This results in a power-law

$$N(\geq E_n) = N_0 \sum_{m=n}^{\infty} (1 - P_{\mathrm{esc}})^m = N_0 \frac{(1-P_{\mathrm{esc}})^n}{P_{\mathrm{esc}}} = \frac{N_0}{P_{\mathrm{esc}}} \left(\frac{E_n}{E_0} \right)^{1-\gamma} , \tag{3.1}$$

with a spectral index $1-\gamma = \ln(1-P_{\mathrm{esc}})/\ln(1+\xi) \simeq -P_{\mathrm{esc}}/\xi$ and the spectrum $dN/dE \propto E^{-\gamma}$. For non-relativistic shocks, this ratio depends only on the compression ratio of the plasma, which gives $\gamma = 2$ for a mono-atomic gas [15, 170]. Relativistic shocks lead to a high energy gain $\Delta E/E$ of order unity. However due to relativistic beaming, only a fraction of particles escapes upstream within a cone of angle $\theta \approx 1/\Gamma_{\mathrm{sh}}$ making the acceleration less efficient. Numerical simulations for this case find slightly softer spectral indices $\gamma \approx 2.2 - 2.3$ [171].

This is, however, only the spectrum that is accelerated inside the source. Depending on the mechanism by which particles escape from the source, it can be significantly modified: Diffusive or direct escape harden the spectrum by up to one power [172], while advective escape can act as a low-pass filter suppressing the high-energy emission [117]. More detailed simulations have found bell-shaped escape spectra, where only the most energetic particles escape [122, 173].

Fermi acceleration can be used to explain UHECR acceleration in a wide range of different source environments, including supernovae remnants (SNR) [174, 175], jets of AGNs [119, 176] and GRBs [177–179], starburst galaxies [127] and galaxy clusters [180, 181]. To account for the observed composition of UHECRs, the source environment needs to be enriched in heavier elements. For AGNs, so-called ‘espresso’ acceleration has been proposed, in which a seed component with Galactic composition enters the jet from the side and is then accelerated to ultra-high-energies [119]. For long GRBs, heavier nuclei from the surrounding stellar medium can be accelerated together with the jet [179].

The main alternative mechanism is direct acceleration, in which cosmic ray particle are accelerated in a strong, extended magnetic field. This idea goes back to Swann [182], who proposed it for stellar magnetic fields in sunspots. For most sources, such electric fields in astrophysical plasmas are not stable, since they are quickly neutralised by electric currents. However, the rapid rotation of highly magnetised compact objects can create a strong electric field through unipolar induction. Pulsars (rotating neutron stars) can produce strong enough fields for UHECR acceleration [126, 183, 184]. The maximal energy of the accelerated cosmic rays depends on the rotational speed of the pulsar. When it spins down, the maximal energy decreases resulting in a hard power-law spectrum $dN/dE \propto E^{-1}$ [183]. Neutron stars are good candidates for explaining the UHECR composition, since their surface is rich in iron. If pure iron is initially extracted from the surface, it will partly disintegrate during acceleration which can produce a mixed composition compatible with the observed UHECR composition [185].

Independent of the exact acceleration mechanism, the maximal cosmic-ray energy is limited by requiring that the particle must be contained inside the acceleration region. This is expressed in the so-called ‘Hillas criterion’ [186], for relativistic shocks it reads [171]:

$$E_{\mathrm{max}} \leq ZeB\beta_{\mathrm{sh}}\Gamma_{\mathrm{sh}}R\,, \tag{3.2}$$

where Ze is the particle charge, B the magnetic field strength in the acceleration region, R the (comoving) radial size of the accelerator, and β_{sh} and Γ_{sh} the shock speed and Lorentz factor. This criterion is typically visualized in a so-called ‘Hillas plot’, which is shown in Fig. 3.1. The plot shows the parameter space in R and B for different selected source candidates. For heavier nuclei, the Hillas criterion is slightly less constraining as they are more efficiently accelerated due to their charge Z. The minimal requirement for the product of $B \cdot R$ to reach $E_{\mathrm{max}} = 10^{20}$ eV is shown as diagonal for protons (red) and iron nuclei (blue) for different shock speeds β_{sh}.

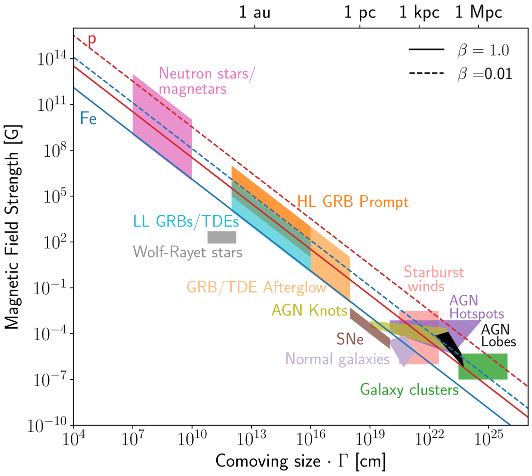

Figure 3.1: Hillas plot: The different source classes are shown as a function of their characteristic size $\Gamma_{\text{sh}} \cdot R$ and magnetic field B, where the values of B are given in the comoving frame. The solid and dashed lines show the limit on the product $\Gamma_{\text{sh}} \cdot R \cdot B$ above which protons (red) or iron (blue) can be confined inside the source depending on β_{sh}, see Eq. (3.2). Reprinted from Alves Batista et al. [16] under CC BY 4.0 license.

However, the Hillas criterion is only a necessary condition for the acceleration of UHECRs. Possible source candidates can additionally be constrained by the following criteria [187]:

Luminosity: The sources should be luminous enough to provide the energy for the acceleration mechanism. In most sources UHECR acceleration is accompanied by the emission of gamma-rays. The UHECR luminosity is usually assumed to be proportional to the gamma-ray luminosity, where their ratio is called baryonic loading.

Emissivity: The total number density of sources and their (effective) UHECR luminosity need to be high enough to sustain the observed UHECR flux. The required local injection rate can be estimated as $\sim 10^{44}$ erg Mpc yr^{-1} [188].

Radiation and interaction losses: Particles can lose energy both by radiation processes in the accelerating field (synchrotron) and due to interactions (photo-hadronic etc.). These processes limit the obtainable maximal energy, if their loss rates exceed the acceleration rate $t_{\text{acc}}^{-1} = \eta c / R_L$ (with the Larmor radius $R_L = E/ZeB$ and acceleration efficiency $\eta \leq 1$).

Secondary radiation: Interactions inside the source will produce secondaries depending on the radiation densities. These secondaries, like neutrinos and photons, should not exceed the limits from observations, both for a single observed source and for the diffuse fluxes.

Radiation and interaction losses depend both on the radiation density in the source and the shape of the target photon spectra. These can be very different for different source types. For example GRB photon spectra can typically be described by a simple broken power law, while AGN photon spectra have more complicated shapes. The energy losses, disintegration and secondary particle production therefore have to be discussed for each source type independently, see e.g. Rodrigues et al. [118] for AGN and Biehl et al. [124] for GRBs. Details on the GRB specific physics are discussed in the context of multi-collision models in Chapter 6. Most of the relevant radiation processes also take place during extragalactic propagation, even though the photon fields have much lower energies than those inside sources. We will therefore discuss these interactions in the UHECR propagation context in Section 3.2.

If the energy losses play a subdominant then the acceleration is only limited by the Hillas criterion. This results in a so-called 'Peter's cycle' [189], where nuclei are accelerated to energies Z times higher than protons. This is specifically the case for sources with low luminosities. This case naturally fits the trend in the composition observed by Auger, which gets increasingly heavy with energies above the *ankle*. It is therefore commonly used as a generalised model of the source spectrum for UHECR propagation [27, 148, 151].

3.2 Propagation of UHECRs

3.2.1 Extragalactic photon fields

During extragalactic propagation, UHECRs lose energy by interacting with extragalactic background photons. The relevant backgrounds are the Cosmic Microwave Background (CMB) and the Extragalactic Background Light (EBL). All diffuse radiation from star forming processes is summarized as the EBL. As most of the EBL falls into the infrared range, the term Cosmic Infrared Background (CIB) is sometimes used synonymously. For a recent review of the measurements of the photon backgrounds in different wavelengths, see [190]. The redshift scaling of the EBL and CMB is displayed in Fig. 3.2.

The CMB was first discovered by Penzias and Wilson [191] in 1964. Today the CMB is explained by standard cosmology as radiation from the time of recombination, constituting one of the main evidences for the Big Bang theory. The CMB is an almost perfect black body spectrum with a current day temperature of $T_0 = 2.726$ K, which corresponds to an average photon energy of $6.34 \cdot 10^{-4}$ eV [192]. It does not receive additional injection and therefore scales only adiabatically. Its redshift dependence is described by simply scaling the temperature as $T = (1 + z)\, T_0$. Due to its high number density (400 photons per cm^3), the CMB is the main target for UHECR interactions.

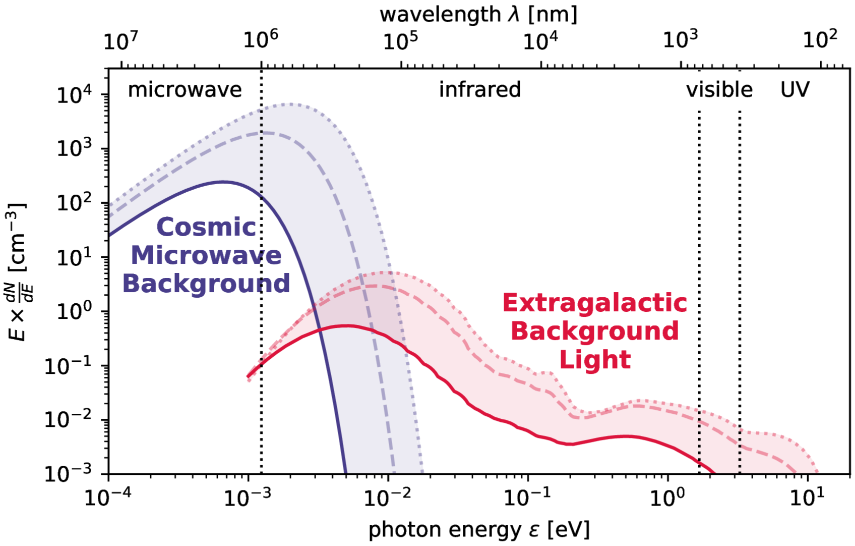

Figure 3.2: Extragalactic photon backgrounds at redshifts between 0 and 2. The Cosmic Microwave Background (CMB, purple) is shown together with the Extragalactic Background Light (EBL, red) at redshifts 0 (solid), 1 (dashed) and 2 (dotted). The CMB originates from the epoch of recombination and therefore scales purely adiabatically. The EBL receives constant injection from sources and therefore scales less strongly with redshift. The scaling with redshift directly impacts the scaling of the interactions rates in Fig. 3.3.

The EBL consists of the cumulative, redshifted radiation from stars and galaxies. Its number density is about 200 times smaller than that of the CMB. However, UHECR interactions with the EBL produce secondaries at lower energies. The EBL therefore has a significant impact on the spectral shape of cosmogenic neutrinos. The EBL has two peaks: The direct contribution of starlight peaks around 1 eV, while a significant fraction is absorbed by interstellar dust and remitted in thermal radiation, producing a second peak around 10^{-2} eV. The EBL receives constant injection following the star formation rate and therefore has a non-trivial redshift evolution. Direct measurements are complicated due to the large backgrounds from local sources. Models rely on different techniques to infer the EBL redshift evolution: Semi-analytical models estimate the EBL by forward evolution from models of the formation of stars and galaxies [193–195] or by using estimates of the local EBL from measurements together with backward evolution models [196, 197]. Empirical models rely on estimating the gamma-ray opacity of the Universe due to the EBL from the spectra of source catalogues [198, 199]. The model by Gilmore et al. [194] is shown in Fig. 3.2 to illustrate the non-linear redshift dependence.

3.2.2 Interactions of UHECRs in extragalactic space

The extragalactic photon environments are homogenous and isotropic on large scales. The same assumption can be made about the distribution of many classes of astrophysical sources. Under this assumption, UHECR propagation can be expressed as a one-dimensional differential equation in energy. The injection from sources is expressed by a single energy- and time-dependent injection function that can either be the output of a source model or a generic analytic form. The arising coupled differential equation system then reads for the homogenous and isotropic particle density $n = \frac{\partial N}{\partial E \partial V}$

$$\partial_t n_i(E,t) = -\partial_E \left(\frac{dE}{dt} n_i(E,t) \right) - \Gamma_i n_i(E,t) + \sum_j Q_{j \to i}(n_j(E,t)) + J_i(E,t) \,. \tag{3.3}$$

The index i denotes a single species, making this a coupled differential equation system. The first term denotes continuous energy losses with an energy loss rate $\frac{dE}{dt}$. The second and third terms denote interaction (Γ_i) and subsequent reinjection $Q_{j\to i}$ at a different energy. Reinjected particles can be either of the same or a different species, many channels produce multiple secondary species. Channels that reinject the same species with slightly changed energy can be approximated by a continuous energy loss term, as we will discuss at the example of photo-pair production. Lastly J_i denotes the external injection of particles. See Appendix A.1 for details on the individual terms of Eq. (3.3).

To account for magnetic diffusion, additional terms would have to be introduced of the form $\mathrm{div}(D(E,x,t)\nabla n(E,x,t))$ with a diffusion coefficient D. This makes the differential equation system significantly more complicated to solve, as it introduces (at least) one spatial dimension x. At lower energies, magnetic diffusion increases the time of propagation significantly, which can become larger than the age of the Universe. This leads to an effective magnetic horizon for UHECRs with energies below $\sim 10^{18}$ eV, i.e. below the *ankle* [161, 200]. Magnetic diffusion is negligible if the diffusion length is larger than the separation between neighboring sources. This can be assumed to a good approximation above the *ankle* [201]. Within this thesis, we therefore neglect the effect of magnetic diffusion.

For a photon field n_γ, the general form of the interaction rate on a target photon field is given by an integral over the photon energy ε and the pitch angle θ

$$\Gamma_i(E_i) = \int \mathrm{d}\varepsilon \int_{-1}^{+1} \frac{\mathrm{d}\cos\theta}{2} (1-\cos\theta) n_\gamma(\varepsilon, \cos\theta)\, \sigma_i(\varepsilon_r) \,, \tag{3.4}$$

where the interaction cross section is typically known in terms of the photon energy in the particle rest frame $\varepsilon_r = (E_i \varepsilon)/m_i$. The reinjection rate has an analogous form, with the differential

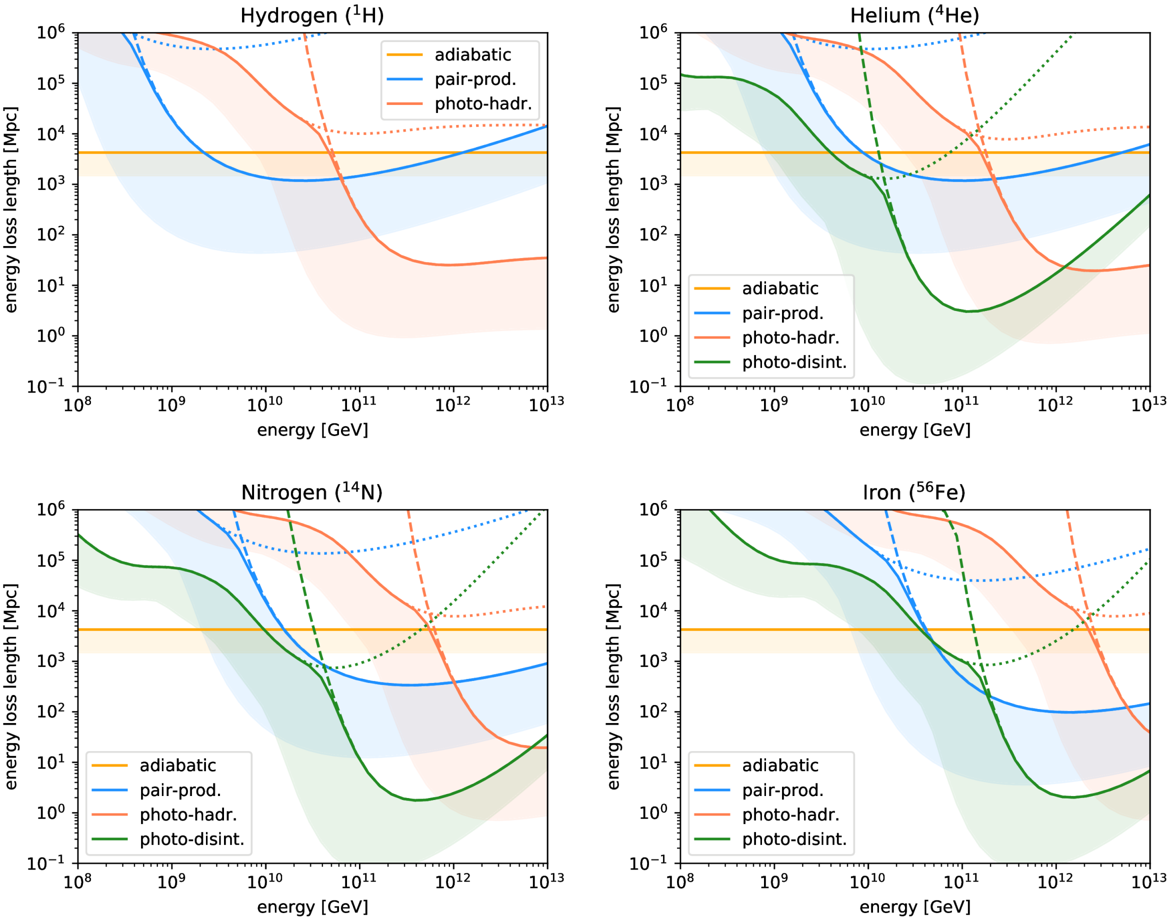

Figure 3.3: Energy loss length in Mpc for protons, helium, nitrogen and iron as a function of energy. Energy loss length are given for adiabatic cooling (see Section 3.2.3), photo-pair-production (see Section 3.2.4), photo-pion-production (see Section 3.2.5) and photo-disintegration (see Section 3.2.4). The rates at local redshift are given by solid lines, with the individual contributions from CMB (dashed) and EBL (dotted) marked individually. The scaling up to redshift 2 is indicated by the shaded region. The redshift scaling of these interaction rates is due to the scaling of the photon fields, which is displayed in Fig. 3.2.

cross section $d\sigma_{j\to i}/dE_i$. The relevant interactions for extragalactic propagation of UHECRs are photo-pair-production, photo-pion-production and photo-disintegration. Furthermore, all relativistic particles are subject to adiabatic cooling due to the cosmological expansion of the Universe. The energy loss lengths for these processes are displayed in Fig. 3.3. All of these processes will be explained in more detail in the following. The processes that dominate at different energy ranges are however already visible in Fig. 3.3. The highest energies are dominated by photo-pion-production (for protons) or photo-disintegration (for nuclei). Pair-production is only relevant for protons, for nuclei it is overshadowed by disintegration. Adiabatic cooling has the same small loss length independent of energy and therefore becomes relevant only below the threshold of the other processes.

Note that source models often assume a single isotropic radiation zone. The arising differential equation system then has essentially the same form as Eq. (3.3), with a few added processes such as e.g. particle escape and synchrotron radiation.

3.2.3 Cosmological expansion of the Universe

As UHECRs are highly relativistic particles and travel over cosmological distances, they are affected by the cosmological expansion of the Universe. This expansion is typically expressed by the scale factor a or redshift z, which are related by $a = (1+z)^{-1}$. a scales the proper distance between two points in spacetime as a function of redshift. The redshift evolution of a is given by the Hubble parameter $H = \dot{a}/a$:

$$H(z) = H_0 \sqrt{\Omega_m (1+z)^3 + \Omega_\Lambda}\,, \tag{3.5}$$

with the Hubble parameter $H_0 = 67.8$ km s^{-1} Mpc^{-1}, the pressureless matter density $\Omega_m = 0.308 \pm 0.012$ and dark energy density $\Omega_\Lambda = 0.692 \pm 0.012$ [192]. Due to this expansion, all relativistic particles emitted at redshift z are observed with an energy E, which is redshifted compared to their initial energy $E(z) = E_0(1+z)$. The local rate of this ‘adiabatic’ energy loss can be expressed as:

$$\frac{dE}{dt}(z) = -EH(z) \tag{3.6}$$

This energy loss is relatively small, however it still becomes relevant for cosmic rays at energies below 2 EeV, where interactions with the photon-backgrounds are negligible. At high energies, the cosmological expansion still plays an indirect role as the number of particles per proper volume decreases as $(1+z)^{-3}$. This is typically accounted for by expressing Eq. (3.3) in comoving densities ($Y = a^3 \cdot n$). As the photon-backgrounds are also subject to this evolution, they were more dense and had higher energies in the past. This both shifts the thresholds for interactions to lower energies and increases their amplitude as indicated by the grey shading in Fig. 3.3.

3.2.4 Photo-pair production

The interaction with the lowest energy threshold for UHECR protons is the creation of an electron-positron pair in the interaction with background photons (Beithe-Heitler process [202]):

$$N + \gamma \rightarrow N + e^{+} + e^{-}. \tag{3.7}$$

This process has a threshold of $\varepsilon_r = 2m_e c^2$, which corresponds to a UHECR proton energy of $1.6 \cdot 10^9$ GeV for CMB photons. Additionally, it has the highest cross section out of all relevant interactions in extragalactic propagation. However, the energy loss per interaction is very small with $\sim 1\%$ per interaction. The energy loss rate is therefore comparably small and pair-production is only relevant below the threshold of photo-hadronic processes. Due to the small energy loss per interaction, pair production losses are typically treated by a continuous energy loss approximation. This loss term was derived analytically by Blumenthal [203] as:

$$\left(\frac{dE}{dt}\right)_{e^+e^-} = -\,\alpha r_0^2 Z^2\, m_e^2 \int_2^\infty d\xi\, n_\gamma \left(\frac{\xi m_e}{2\Gamma}, z\right) \frac{\Phi(\xi)}{\xi^2}\,, \tag{3.8}$$

with the fine-structure constant α and the classical electron radius r_0. The dimensionless function $\Phi(\xi)$ contains the angular averaged cross section and inelasticity. For nuclei this scales as Z^2/A, as the cross section increases with Z^2, while the inelasticity scales as A^{-1}. However, also the threshold increases proportional to A (appearing in the above formula implicitly in the Lorentz factor Γ). Effectively, this threshold scaling makes pair production less relevant in the context of nuclei, their interactions are instead dominated by disintegration.

For protons, the energy loss rate peaks around energies of $5{\cdot}10^9$ GeV$-1{\cdot}10^{10}$ GeV depending on redshift as seen in the upper left panel of Fig. 3.3. This could in principle cause the *ankle* in the observed UHECR spectrum, if it was dominated by protons [138–141]. Still several arguments have since disfavoured this so-called *dip-model*, such as the heavier composition observed by Auger [93] and the very high cosmogenic neutrino and gamma-ray fluxes expected in this case [5, 143, 144].

3.2.5 Photo-nuclear interactions

There are multiple photo-nuclear processes that all lead to the disintegration of the nucleus into two or more fragments; an example is α-emission:

$${}^{A}_{Z}N + \gamma \rightarrow {}^{A}_{Z}N^{*} \rightarrow {}^{A-4}_{Z-2}N + {}^{4}\mathrm{He} \tag{3.9}$$

The lowest photo-disintegration energies are dominated by the ‘giant dipole resonance’ (GDR). In the GDR range ($\varepsilon_r = 8 - 40$ MeV) the photon interacts with the bulk nucleus, which leads to a collective resonant excitation. The subsequent de-excitation leads to the ejection of one

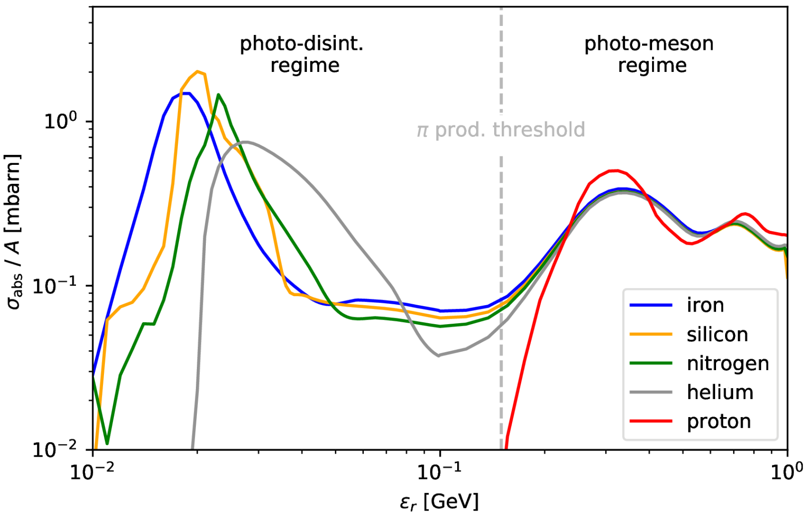

Figure 3.4: Cross section per nucleon by energy regimes for different nuclei (from the Peanut model [204, 205]). In the disintegration regime at lower energies, the interaction is dominated by the Giant Dipole Resonance (GDR). The peak of the GDR shifts to lower energies for higher nucleus mass. Above $1.5 \cdot 10^{-1}$ GeV the pion production regime sets in, which is dominated by baryonic resonances. The cross section in this regime scales directly with the number of nucleons. However, for protons the cross section in this regime is significantly more peaked than for other nuclei, for which it smears out due to Fermi motion of the particles inside the nucleus.

or two nucleons. The GDR peak scales to lower energy for heavier nuclei. At higher energy, 'quasi-deuteron'-processes (QD) take over, in which the photon initially interacts with a pair of nucleons. This is followed by the emission of nucleons or light nuclei. Above energies of 150 GeV, baryonic resonances dominate, which can produce one or - at higher energies - multiple mesons. In the context of UHECR interactions, the first two processes are typically summarized as 'photo-disintegration', while the last is treated separately as 'photo-meson production'. Though, this splitting is somewhat artificial, since photo-meson production also disintegrates the nucleus. The cross section for different elements in these regimes is shown in Fig. 3.4.

Photo-disintegration

There are several cross section measurements for the photo-disintegration of different nuclei; the most complete database is EXFOR [206]. The data, however, is sparse and mainly available for stable nuclei. Additionally, there are no measurements for different isobars, i.e. two or more elements with the same number of nucleons. This constitutes significant uncertainty in the

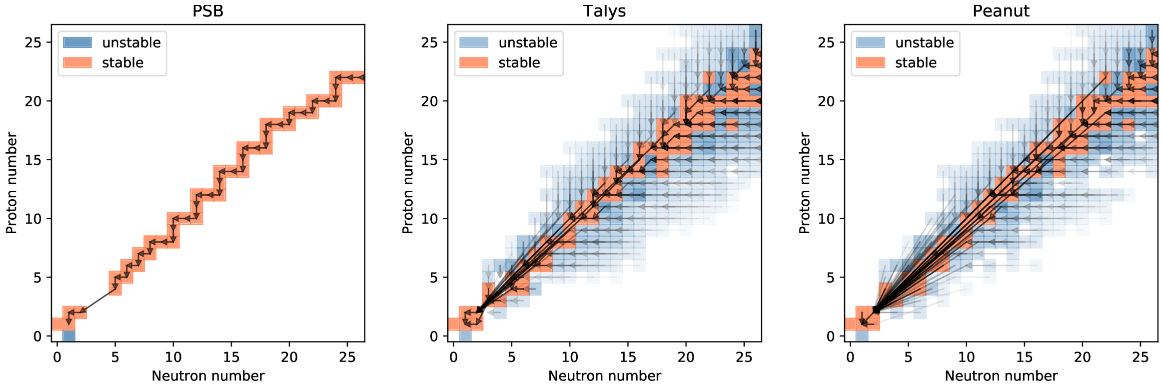

Figure 3.5: Disintegration chains for the disintegration models PSB [207], Talys [208] Peanut [204, 205]. Each isotope is indicated by a square, for unstable nuclei the shading is indicative of there lifetime. For each isotope the main emission channel is shown by an arrow. This channel is defined by the element (excluding hydrogen) emitted with the highest multiplicity around the GDR ($\varepsilon_r \geq 40$ MeV).

astrophysical context, where the chain of disintegration produces a mix of all nuclei including unstable ones. In astrophysical environments, interactions of unstable nuclei become important due to their relativistically dilated lifetime. This is mainly relevant in source environments, while during propagation the time scales are large enough that they mostly decay. The current state of the cross section measurements has been discussed in detail in Boncioli et al. [18].

Due to the insufficient data, a full simulation of the nuclear cascade has to rely on nuclear models for the cross sections of all stable and unstable nuclei. Available models include the Puget-Stecker-Bredekamp model (PSB) [207], Talys [208], and Peanut [204, 205], which is part of the FLUKA code. The Talys model can only be applied to particles with mass number $A \geq 12$. It therefore needs to be complemented by other models for lighter elements. Cross sections for light elements were compiled from different sources for the CRPropa code [209]. The different disintegration models result in qualitative and quantitative differences in the disintegration chain, which is illustrated in Fig. 3.5. Boncioli et al. [18] discussed the comparison of these models to data in detail as well as their impact on astrophysical source simulations.

PSB is an empirical model. It choses a single isotope per mass number and simple parameterizations of the cross sections empirically estimated from data. The unique disintegration chain is populated by subsequent emission of nucleons. Both Talys and Peanut are more sophisticated statistical models. They provide cross sections for unstable isotopes and include disintegration channels into heavier fragments (D, T, ^{3}He, ^{4}He). Compared to PSB this leads to a faster disintegration into light elements, altering the interpretation of the final observed composition. The disintegration chains between Talys and Peanut are qualitatively similar.

There are however quantitative differences in the cross sections depending on the detail of the models and the level of optimization to data. For example Talys treats the total absorption cross section for different isobars the same, while Peanut shows significant differences. Overall, Peanut seems to be better optimized to data [18]. In some channels, Talys shows significant deviation from data; unofficial tables are available to correct for this [210]. PSB follows a simple chain, emitting only a single nucleon per interaction. The production of helium is therefore limited in this model. In contrast, Talys and Peanut are more efficient in helium emission.

Photo-meson production

Above a threshold photon energy of 140 MeV in the CR rest frame, photo-hadronic interactions can excite baryonic resonances in individual nucleons, which results in the production of mesons. At threshold, pion-production is approximated by the Δ^+-resonance:

$$p + \gamma \rightarrow \Delta^+ \rightarrow \begin{cases} \pi^0 + p & 2/3 \text{ of all cases} \\ \pi^+ + n & 1/3 \text{ of all cases} \end{cases} . \tag{3.10}$$

In each interaction, the proton loses a significant fraction of its energy which is determined by the mass ration m_π/m_p (on average 20% around the Δ^+-resonance). These high energy losses per interaction are the reason that pion-production has a much higher energy loss rate compared to photo-pair production, even though its cross section is two orders of magnitude lower. Several other production channels can occur at higher energies, including the production of multiple pions. Including all relevant channels significantly changes the production ratio of π^+ to π^0 [211].

The pions are unstable and decay dominantly as $\pi^+ \rightarrow \mu^+\nu_\mu \rightarrow e^+ + \nu_e + \nu_\mu$ and $\pi^0 \rightarrow \gamma\gamma$. Similarly, the neutron in the charged pion channel will decay as $n \rightarrow p + e^- + \bar{\nu}_e$. The charged pion channel is the main source of astrophysical neutrinos. The expected flavour ratio of neutrinos from photo-pion production is therefore $\nu_e : \nu_\mu : \nu_\tau = 2 : 1 : 0$. In very dense source environments the pion and neutrino can re-interact before decaying or in case of the charged pion be damped by magnetic fields, reducing the flux of neutrinos and changing their flavour ratio. Neutrinos produced during extragalactic propagation are referred to as 'cosmogenic' neutrinos.

Out of the disintegration models discussed in this section, only Peanut reaches into the pion-production regime. This regime is therefore typically covered by a separate model. A widely used model is SOPHIA [212], a Monte-Carlo generator that was initially developed for interactions of free protons and neutrons with photon fields. SOPHIA includes additional hadronic resonances, direct pion production, and the production of multiple pions at higher energies. This gives a more accurate prediction of the neutrino production than the simple Δ^+-

resonance approximation. For nuclei, a superposition model is typically assumed, in which they are modelled as A free nucleons [211, 213, 214]. In these superposition models, the interacting nucleon is then ejected while the residual nucleus is kept intact. Similar to pair-production, this superposition model increases the amplitude of the cross section by the mass number A, but also shifts the threshold to higher energy. For the same maximal energy, UHECR nuclei therefore produce significantly less cosmogenic neutrinos.

A more sophisticated photo-meson production model has recently been presented by Morejon et al. [215]. It uses SOPHIA as a starting point, but makes several modifications. Firstly, the photo-meson production cross sections are modified for nuclei: At lower energies, measurements show that the Δ^+-resonance is broadened by Fermi-motion for nuclei, leading to a universal curve scaling with A. This effect is visible in the cross sections of the PEANUT-model in Fig. 3.4. Furthermore, the pion-production around the Δ^+-resonance scales with $A^{2/3}$ rather than A, since pions can reinteract inside the nucleus. At higher energies, nuclear shadowing effects lead a reduction of the cross section compared to the simple scaling with A. Secondly, the disintegration of the residual nucleus is taken into account. This leads to faster disintegration in the photo-meson regime. Morejon et al. [215] have studied the effects of this model in different source environments. Depending on the shape of the photon fields, this model leads to some difference in the ejected UHECR spectra and neutrino production. As interactions during extragalactic propagation are dominated by the peaked CMB, they happen mainly in the photo-disintegration regime, such that the more sophisticated photo-meson model should have rather little impact. For the work presented in this thesis, we used the simpler SOPHIA-superposition model. Still, some effect from the new model might be visible in the neutrino prodcution from EBL interactions, which might be interesting to test in future work.

After the detection of the CMB, Greisen [136] and Zatsepin and Kuzmin [137] predicted that photo-pion production should lead to a strong suppression of the flux of UHECR protons due to its large energy loss of $\geq 20\%$. This is the so-called *GZK-effect*. As seen in Fig. 3.4, the threshold energy of the GDR is lower for heavier nuclei than that of photo-pion production for protons. This threshold is expressed in terms of the rest frame energy which is de-boosted by the Lorentz factor $\gamma = E/m$. As heavier nuclei have a smaller Lorentz factor for the same kinetic energy, the disintegration threshold in terms of the observed energy E is pushed to higher values for heavier masses. As a consequence the disintegration on the CMB leads to a cutoff in the spectrum of UHECR-nuclei at energies similar to the threshold for photo-pion production for protons. The term *GZK-effect* is therefore sometimes also used in the context of nuclei, though the original publications discussed only protons. Lighter nuclei such as helium disintegrate at energies about an order of magnitude lower. This is illustrated in terms of the interactions rates in Fig. 3.3.

3.3 Hadronic interactions and air showers

As discussed in Section 2.2, the measurement of UHECR composition is based on the indirect observation of extensive air showers. It is therefore important to understand the model assumptions going into the interpretation of the observations. For a more detailed overview see Kampert and Unger [20].

For a primary UHECR baryon, the development of the air shower is driven by QCD interactions with the air atoms. The initial and subsequent interactions produce secondary hadrons, mostly pions, kaons and nucleons. Neutral pions decay almost immediately into photons ($\pi^0 \rightarrow \gamma\gamma$), which cascade electromagnetically trough e^+e^--pair production and bremsstrahlung. At high energies, charged pions and kaons can re-interact before decaying, feeding the hadronic component of the shower. In each particle generation, the hadronic component transfers $\sim 25\%$ of its energy to the electromagnetic component. At lower energies, mesons have time to decay producing muons and neutrinos ($K^{\pm}/\pi^{\pm} \rightarrow \mu^{\pm} + \nu_{\mu}/\bar{\nu}_{\mu}$). Eventually, about 90% of the shower energy is dissipated in the electromagnetic component, while the rest ends up in the muonic component [92].

The longitudinal development of the number of shower particles can be modelled by the 'Gaisser-Hillas' function [216]:

$$N(X) = N_{\max} \left(\frac{X - X_0}{X_{\max} - X_0} \right)^{\frac{X_{\max} - X_0}{\lambda}} \exp\left(\frac{X_{\max} - X}{\lambda} \right) \tag{3.11}$$

with the initial interaction at depth X_0 and the depth of the shower maximum $X_{\max}$, the number of particles at the shower maximum $N_{\max}$ and a shape parameter λ. In principle, the shape parameter can be used as a composition related observable, however this is not possible with current systematic uncertainties [217]. Another observable that is in principle related to composition is the muon content of the air shower. Measurements from both TA and Auger however seem to show an excess of muons over the model predictions of around 50% [218, 219]. Muon measurements are therefore currently not conclusive for composition measurements. The composition is therefore mostly inferred from the average distribution of $X_{\max}$, or rather its first two moments $\langle X_{\max} \rangle$ and $\sigma(X_{\max})$. In the future, improvements of the models as well as reduced systematics of the upgraded experiments, will allow to take additional observables into account. This might enable a measurement of the composition of individual showers.

At lower energies, both electromagnetic and hadronic interactions are well understood. However for UHECRs, the initial particles as well as some of the early secondaries interact at energies which are well above the energy scale of the LHC. The shower development is very sensitive to the cross sections and distribution of secondaries in these initial interactions. Several phenomenological models try to extrapolate the hadronic interactions to higher energies. Among the ones tuned to LHC data are Epos-LHC [220], Sibyll 2.3 [221] and QGSjet04-II [222].

For now, the main composition related observable remains the average distribution of $X_{\max}$ between showers. It is sensitive to composition since showers of heavier elements develop earlier in the atmosphere than protons. Qualitatively, this can be understood by viewing them as A free nucleons. The shower then consists of A particles in the initial interaction, compared to a single particle in the proton case and therefore develops earlier. A more detailed understanding requires full shower simulations. A widely used Monte-Carlo code for these simulations is CORSIKA [223]. The dependence of the composition related observables $\langle X_{\max}\rangle$ and $\sigma(X_{\max})$ has been analysed in Abreu et al. [19] for CORSIKA simulations with different hadronic interaction models. They found that the energy dependence is well described by the empirical formula:

$$\langle X_{\max}\rangle(E) = \langle X_{\max}\rangle_p(E) + f_E(E)\langle \ln A\rangle \tag{3.12}$$

where $\langle X_{\max}\rangle_p(E)$ is the (model-dependent) depth of shower maximum for protons and $\langle X_{\max}\rangle$ scales linearly with $\langle \ln A\rangle$ with the energy dependent coefficient $f(E)$. The variance is given by

$$\sigma^2\left(X_{\max}\right)(E) = \sigma^2_{\mathrm{sh}}(\langle \ln A\rangle) + \left[b\,\sigma_p^2(E) + f_E^2(E)\right]\sigma^2(\ln A) \tag{3.13}$$

where $\sigma^2_{\mathrm{sh}}(\ln A) = \sigma_p[1 + a\ln A + b(\ln A)^2]$ describes the shower-to-shower fluctuations for pure composition and $\sigma_p(E)$ is the fluctuation for pure proton composition. Only the second term depends on the variance of masses, while the first scales with the average mass. These formulas are only empirical and contain several free parameters (a, b and implicitly in $\langle X_{\max}\rangle_p$, σ_p and $f(E)$) that need to be determined by shower simulations. We use an updated set of parameters for the post LHC models in this thesis[1], compared to the ones provided in Abreu et al. [19].

The average $\langle X_{\max}\rangle$ scales only with the average mass. However, the second moment $\sigma(X_{\max})$ depends both on the average mass as well as the variance in masses. The shower-to-shower fluctuation will be highest for a mix of different masses. The observed $\sigma(X_{\max})$ is therefore an indicator of how pure the observed composition is. The connection between the mean and variance of $X_{\max}$ becomes more clear in so-called *umbrella* plots [20], which plot the $X_{\max}$ data in $\sigma(X_{\max})$ over $\langle X_{\max}\rangle$. Such a plot is displayed in Fig. 3.6. It shows simulations for different linear transitions between the masses ($H \to Fe$, $Fe \to Si$, $Si \to N$, $N \to He$ and $He \to H$) in each energy bin. The energy dependence is mostly absorbed by subtracting the values for pure iron from both the data and simulations in each energy bin. We choose here to subtract the result from Epos-LHC. For any mix of elements, the average mass must always be between

---

[1] Private communication with S. Petrera

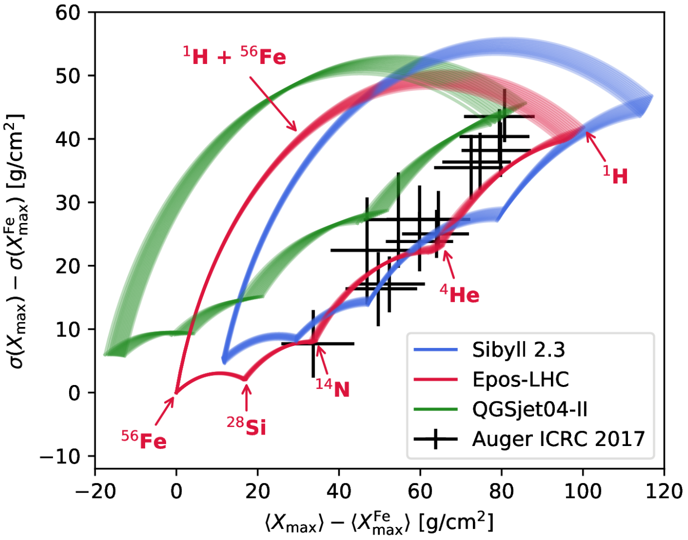

Figure 3.6: *Umbrella* plot of the Auger composition measurements. The Auger composition data [93] is shown in $\langle X_{\rm max}\rangle$ over $\sigma(X_{\rm max})$. Error bars are shown as combined statistical and systematic uncertainties. For each energy bin, the theoretical prediction from the Epos-LHC model is subtracted to filter out the linear energy dependence. Predictions for the models Epos-LHC, Sibyll 2.3 and QGSjet04-II are shown at different energies for linear transitions between different elements (^{1}H, ^{4} He, ^{14}N, ^{28}Si,^{56}Fe). The spread in the curves comes from the non-linear energy dependence. The maximal variance in $\ln(A)$ is given for a mix of ^1H and ^{56}Fe. Any mix of masses must therefore be contained inside the umbrella shape.

iron and protons, while the highest mass dispersion is given for an equal fraction of iron and protons. Therefore, any combination of masses must be contained in the umbrella shape.

From Fig. 3.6, one can infer the impact of the hadronic model. Epos-LHC predicts a light composition at lower energies with a significant abundance of protons. It becomes increasingly heavier with energy up with the last data point, that is consistent with pure nitrogen. The predicted $\sigma(X_{\rm max})$ leaves little room for a superposition of masses. Compared to that, Sibyll 2.3 gives a heavier interpretation of the $X_{\rm max}$ data with the last data point close to silicon. It also leaves slightly more room for a superposition of masses. The only model that cannot produce a consistent interpretation of the observed data is QGSjet04-II. Already for pure composition, it predicts a $\sigma(X_{\rm max})$ that is higher than the observed fluctuations. See Abreu et al. [19] and Bellido [93] for details on this contradiction between QGSjet04-II and Auger data. Considering QGSjet04-II only for $\langle X_{\rm max}\rangle$, it would predict the lightest composition out of the three models.

Chapter 4

PriNCe: a new code for UHECR propagation

During the work leading up to this thesis, new code for the extragalactic propagation of UHECRs was developed from from the ground up in collaboration with Anatoli Fedynitch. The code is called PriNCe[1] (short for ***Pr**opagation **i**ncluding **N**uclear **C**ascade **e**quations*). The basic idea for this code was to explicitly describe the coupled differential equation system in terms of vectorized expressions. This allows for easy numerical optimizations, by using computational libraries optimized for linear algebra. The relevant physical interactions are described in the previous chapter, in Section 3.2. This chapter serves as an introduction to the basic principles used for vectorization in PriNCe and gives an overview of its capabilities and possible future extensions. More details on the methods can be found in Appendix A.

4.1 Idea and motivation

The two UHECR-propagation codes mainly used in the literature, CRPropa [214] and SimProp [213], both use a Monte-Carlo approach. This means they propagate individual particles for different initial parameters like redshift and energy sampled from a distribution. For a large sample of particles, the resulting spectrum then converges towards the exact solution. An advantage of this approach is that the sampling does not have to be recomputed for every new parameter set. The result for different distributions in the initial injection can be approximated by re-weighting each event of the initial sample by its parameters. For example for a changed spectral index γ_{new} of the energy spectrum, each simulated event can simply be re-weighted with the weight $w = E^{-\gamma_{\mathrm{new}}+\gamma_{\mathrm{old}}}$. For accurate results, this re-weighting method needs a large number of events in the initial sample to achieve exact results. While computing the initial sample is costly in computation time, the re-weighting can be executed very fast.

CRPropa accelerates the speed in its calculations by precomputing and storing all interaction rates. If the interaction rates change, the whole event sample has to be recomputed. This is the case for changing photon fields and cross sections. The Monte-Carlo approach therefore

---

[1] The source code is available at `https://github.com/joheinze/PriNCe`

becomes slow when studying systematics of cross sections or non-linear scaling of photon fields. The same also constitutes a problem for the application to source physics, where the photon fields can have a wide range of shapes and time dependencies.

An alternative approach is to directly solve the transport equation for the particle density dN/dE. For protons only, this approach is relatively straightforward to implement, as they produce only a few secondary particle species. However, for the full nuclear cascade, the transport equation becomes a system of ~ 100 coupled integro-differential equations, which requires more effort to implement numerically and greatly increases the required computation time. For comparable numerical optimization, the computation of a single spectrum in this approach is faster than for the Monte-Carlo method. However, this approach does not allow for a re-weighting of precomputed samples. The Monte-Carlo method can therefore be more efficient depending on the parameters of interest. Solving the transport equation directly is specifically advantageous for parameters where re-weighting is not possible.

An example for this method is the NeuCosmA library, that implements the full nuclear cascade in a computation efficient way [124]. Its speed results from a low-level implementation in C. Similar to CRPropa, the interaction rates are computed once at the start of the computation. This approach becomes slow for time dependent photon fields, as they require many recomputations. Therefore, NeuCosmA has, until now, only been applied for the nuclear cascade in sources [118, 130]. An implementation for the propagation of UHECR protons exists only for the proton case [224].

To overcome these limitations, we chose to re-implement the UHECR propagation in a new code. We use the high-level language Python to keep the code accessible and modular. Computational performance is achieved by rewriting the transport equation in terms of vectorized expressions. This allows them to be accelerated by using scientific libraries like numpy and scipy [225]. Details will be given in this chapter. The new propagation code is called PriNCe. To summarize, the main development goals were the following:

Efficient, time-dependent UHECR solver: The code should recompute the interaction rates explicitly in every step. Still, the code should be efficient enough to compute a single spectrum within seconds, to be competitive with optimized Monte-Carlo codes.

Easy variation of model inputs: The code should allow to easily switch out model inputs like photon-fields and cross sections. We therefore do not rely on any pre-computations

Accessibility and modularity: By writing the code in an object-oriented way, it should be easy for users to modify and extent specific parts.

PriNCe is designed to be compatible with the data files and analytic conventions of Neu-

CosmA. However, it was written from the ground up using Python. The basic concept of rewriting the transport equation as a vectorized expression is derived from the MCEq code [226], which solves the transport problem for atmospheric cascades.

4.2 Matrix formulation of the coupled differential equation

The relevant physical processes for the propagation of UHECRs have been introduced in Section 3.2. We show here the basic transport equation for the particle species i (3.3) again:

$$\partial_t Y_i = -\partial_E (b_{\mathrm{ad}} Y_i) - \partial_E \left(b_{e^+e^-} Y_i\right) - \Gamma_i Y_i + \sum_j Q_{j\to i}(Y_j) + J_i. \tag{4.1}$$

The particle indices in this equation go over all system species, $i, j \in (\nu_e, \ldots, \mathrm{p}, \ldots, \mathrm{Fe})$. The individual terms are explained in greater detail in Appendix A.1.

We will omit energy loss terms for adiabatic cooling and pair production for now. Expanding the re-injection term $Q(j \to i)$ the integro-differential nature of the transport equation becomes clear:

$$\partial_t Y_i = -\Gamma_i(E_i) Y_i(E_i) + \sum_j \int_{E_i}^{\infty} dE_j \frac{d\Gamma_{j\to i}}{dE_i}(E_j, E_i) Y_j(E_j) + J_i(E_i). \tag{4.2}$$

In each step one has to integrate over all parent species to get the re-injection for a single particle species i. This problem has no analytical solution. Therefore any implementation has to perform this integral by some numerical integration method. The most simple integration method is step integration. While many computational libraries implement generic integration functions, it is more efficient to explicitly rewrite the equations in a discretized way and implement them on a low computational level. We will now rewrite the transport equation system explicitly in terms step integrals, to show that this leads to vectorized expressions. To solve the problem numerically, it has to be expressed on a grid. We choose a logarithmic grid in energy:

$$E^k = E^0 \cdot 10^{k \cdot d_k} \quad \text{with bin widths} \quad \Delta E_i = E_i - E_{i+1}\,. \tag{4.3}$$

The grid spacing d_k can be adjusted to achieve the required accuracy. For convenience, we describe all quantities defined on this grid in a short notation, e.g. for the particle density $Y_i(E^k) \to Y_i^k$. We found that 8 bins per decade currently produce sufficiently accurate results. It is convenient to use the same grid for all species. However, our approach does not necessarily require this; the grid can be adjusted for individual species, if higher accuracy is required. In the following, we will use k, l, m as upper indices for the energy grid indices and i, j as lower indices for particle species. Quantities like Y_i^k are in most cases represented by their values on the bin center. In some cases, such as for strongly peaked cross sections, it is necessary to compute the precise averages over each energy bin instead of taking the central values.

Expressing the integral in Eq. (4.2) as a step integrals ($\int_{E^i}^{E^j} f(E)dE \rightarrow \sum_{k=i}^{j} f^k \Delta E^k$) the transport equation can be written as:

$$\partial_t Y_i^k = -\Gamma_i^k Y_i^k + \sum_j \sum_{l \geq k} \Delta E_j^l \Delta\Gamma_{ij}^{kl} Y_j^l + J_i^l, \tag{4.4}$$

where $\Delta\Gamma_{ij}^{kl} = \frac{d\Gamma_{j\to i}}{dE_i}(E_j^l, E_i^k)$. This is an ordinary differential equation of degree $N_{\mathrm{particle}} \times N_{\mathrm{E-grid}}$. This system can be rewritten again by defining a single particle state vector as:

$$\vec{Y} = (Y_{\nu_e}^0 \dots Y_{\nu_e}^K \dots Y_p^0 \dots Y_p^K \dots Y_{\mathrm{Fe}}^0 \dots Y_{\mathrm{Fe}}^K)^T\,, \tag{4.5}$$

and an interaction matrix containing the interaction(Γ_i^k) and reinjection ($\Delta\Gamma_{ij}^{kl}$) coefficients as:

$$\Phi_{ij}^{kl} = \begin{cases} -\Gamma_i^k + \Delta E_j^l \Delta\Gamma_{ij}^{kl} & \text{if } i = j \text{ and } k = l \\ \quad\;\; + \Delta E_j^l \Delta\Gamma_{ij}^{kl} & \text{if } i \neq j \text{ or } k \neq l \end{cases}. \tag{4.6}$$

The coupled differential equation system then becomes a simple vectorized expression:

$$\partial_t \vec{Y} = \Phi \times \vec{Y} + \vec{J}, \tag{4.7}$$

where the injection terms in $\vec{J}$ are ordered in the same way as $\vec{Y}$. The interaction matrix Φ is displayed in Fig. 4.1 for $z = 1$. In this case, we assumed boost conservation for the disintegration of nuclei, meaning all fragments are ejected with the same Lorentz factor. For nuclei the energy per nucleon $E_A = E/A$ is conserved. It is therefore advantageous to express the equations in E_A, details can be found in Appendix A.2. Expressed in E_A, all submatrices of Φ coupling nuclear species are diagonal. Protons and leptons are not produced under boost conservation, their submatrices are therefore not diagonal. We will refer to these as 'redistributed' (in energy). Notice that Φ is sparse, i.e. only $\sim 2\%$ of its entries are non-zero. We chose the ordering of $\vec{Y}$, such that it is ordered by mass and energy. Due to this ordering Φ is an upper-triangular matrix as long as there is no particle acceleration. Both of these features of Φ have significant impact on the numerical performance, as we explain in the following section.

Until now, we omitted continuous energy losses and particle decay in the derivation of Eq. (4.7). Expressing the continuous loss terms of the form $\partial_E(bY_i)$ in finite-differences, leads to a matrix expression with a banded matrix for the E-derviative (see see Appendix A.4 for details). Particle decays are completely analogous to interaction and re-injection (see Appendix A.1 for details). Both processes can therefore be absorbed into the interaction matrix Φ.

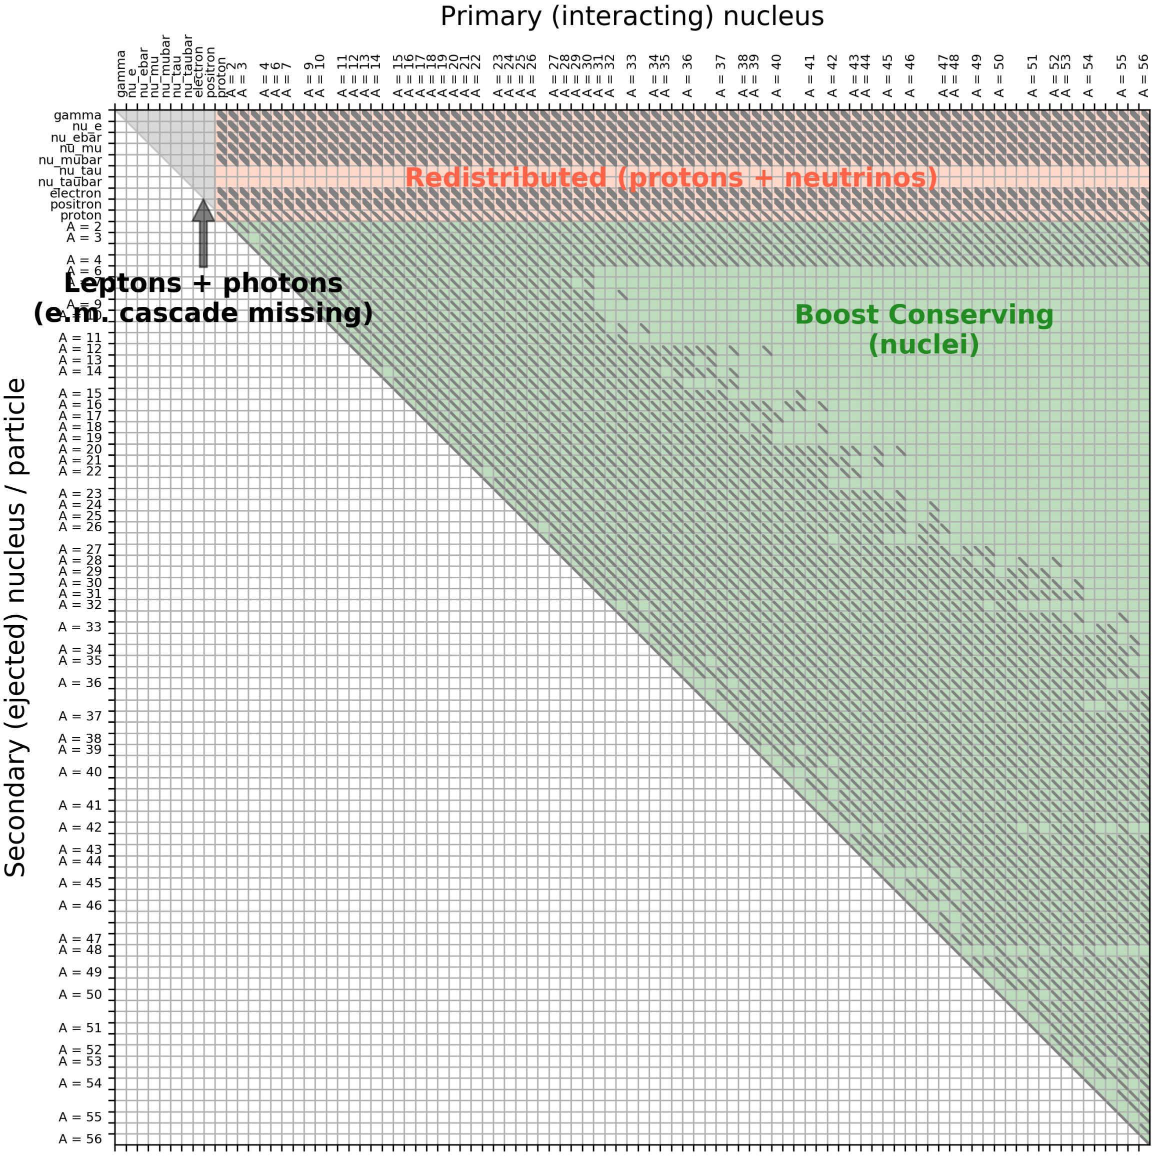

Figure 4.1: Structure of the coupling matrix Φ_{ij}^{kl} for the Peanut model at $z = 1$. The matrix is flattened to two-dimensions $\Phi_{ij}^{kl} \to \Phi_{\alpha\beta}$ by defining a single index for energy and particle index, i.e. $\alpha \equiv (i, k)$, $\beta \equiv (j, l)$ (and the ordering as in Eq. 4.5). The matrix shows the interaction rate of all species (x-axis) and their reinjection into secondaries (y-axis). For leptons, the axes are labelled by names, while for nuclei each mass group is labelled by the mass number A. The gray elements indicate non-zero matrix elements (but not their magnitude). Interaction rates appear on the main diagonal. The reinjection into secondary nuclei are diagonal submatrices, as the energy per nucleon is conserved in our model of disintegration (green part). Secondary leptons and protons show a distribution in energy, leading to triangular submatrices (orange part). There are no interactions for leptons and photons, leading to an empty submatrix (gray part). Implementing the electromagnetic cascade would introduce couplings between electrons, positrons and photons.

4.3 Sparse matrix algebra

As mentioned before, the interaction matrix Φ is sparse. This allows to use sparse matrix algebra, which significantly affects numerical performance. The idea of sparse matrix formats is to not store every element of a matrix, but only the non-zeros ones along with an index scheme that translates them to their location in the matrix. We currently use the SCIPY.SPARSE package, which implements the following formats:

COO *(coordinate format)* Stores the data for each non-zero entry, along with its column and row indices. Efficient for sparse, but unordered matrices.

CSR and CSC *(compact sparse row and compact sparse column)* For CSR, the non-zero data array and the column index array are stored ordered by their rows. A row index pointer stores the location of each row in this array. CSC works analoguely, but with the data ordered by columns. Efficient when there are either a lot of empty row or columns respectively.

BSR *(block sparse row)* Similar to CSR and CSC, but the data array is ordered such that non-zero data blocks are stored together. Efficient when the matrix has many dense (square) blocks.

DIA *(diagonal format)* Stores all (off-)diagonals that contain non-zero elements along with their offsets. Efficient when only few off-diagonals in the matrix are populated.

By looking at the matrix in Fig. 4.1, it is not obvious which format is most suited for the compression. We therefore directly compare computation times in Tab. 4.1 for different parts of the matrix: the full matrix, only nuclear species (including protons), and all nuclei except protons. The CSC format gives the best performance for the full matrix including secondaries. We therefore chose this format in our current implementation. Notice however that if only the nuclear species heavier than protons are included in the system, the DIA format performs best. This means a lot of performance goes into the computation of the secondaries, which are redistributed in energy. By splitting the matrix into submatrices stored in different formats, this could be used for optimizations in the future.

Apart from this direct performance improvement, the sparse matrix structure can be exploited for another performance boost. Notice that all sparse matrix formats save the matrix as a separate data vector with an indexing scheme to translate the data into a dense matrix. For a given cross section model, the interacting particles and their secondaries are fixed. In turn, this means that the location of the non-zero elements in the interaction matrix Φ are also fixed. In

	full matrix		only nuclear species		only boost conserving	
format	size [MB]	speed [ms]	size [MB]	speed [ms]	size [MB]	speed [ms]
CSR	24.3	2.35	4.19	0.33	3.54	0.27
CSC	24.3	1.71	4.19	0.29	3.54	0.25
BSR	21.8	2.57	4.19	0.33	3.54	0.27
COO	32.3	5.13	5.55	0.75	4.69	0.62
DIA	184.00	10.00	38.00	1.67	4.68	0.20
dense	511.00	39.10	417	3100	407.00	30.80

Table 4.1: Computation times for the product $\Phi \times \vec{Y}$ for different sparse matrix formats and dense format (all entries stored explicitly). They are compared for the three cases: the full matrix; accounting only for the ejection of nuclear species, but no neutrinos and photons; for only the boost conserving part (blue part in Fig. 4.1, same as the second case, but no protons)

each integration step, we therefore only have to replace the data vector instead of recomputing the full matrix and determining the sparsity structure. By ordering the cross sections according to the sparse matrix format into a kernel matrix, this can be done in a single matrix expression. See the Appendix A.3 for details. This methods allows to recompute all the interaction rates for a changed photon field in a few milliseconds. While the kernel matrix takes some time to computed it can be stored once for a cross section model. Single cross section can then be rescaled or replaced later with little computational effort.

4.4 Numerical solver and stiffness

Assuming that the interaction matrix Φ and the injection vector $\vec{J}$ are both constant in time, the formal solution for Eq. (4.7) reads:

$$\vec{Y}(t) = \left(e^{\Phi t} - 1\right) \cdot \left(\Phi^{-1}\vec{J}\right) \tag{4.8}$$

or expressed in terms of the eigenvalues λ and eigenvectors $\vec{\Psi}$ of Φ:

$$\vec{Y}(t) = \sum_i \left(e^{\lambda_i t} - 1\right) \cdot \left(\frac{C_i}{\lambda_i}\right) \vec{\Psi}_i \,, \tag{4.9}$$

where $\vec{J} = \sum_i C_i \vec{\Psi}_i$ This means the system converges exponentially with t against the steady state $(-1) \cdot (\Phi^{-1}\vec{J})$, as long as the eigenvalues λ are smaller than 0. However as discussed before, both Φ and $\vec{J}$ are explicitly time dependent for UHECR propagation. We therefore need to treat the system by numerical integration.[2] Using explicit Euler integration a time step ($t_{n+1} = t_n + \Delta t$) becomes:

$$\vec{Y}_{n+1} = \vec{Y}_n + \Phi(t_n) \times \vec{Y}_n \Delta t + \vec{J}(t_n)\Delta t \tag{4.10}$$

$$= (1 + \Phi\Delta t) \times \vec{Y}_n + \vec{J}\Delta t \tag{4.11}$$

If we require $\vec{Y}(t)$ to converge, this implies that $|1 + \lambda_i \Delta t|$ must be smaller than 1 for all eigenvalues of Φ. This leads to the stability criterion (called A-stability):

$$\Delta t = \frac{dt}{dz}\Delta z \leq \frac{2}{\lambda_i} \,. \tag{4.12}$$

The distribution of eigenvalues of Φ at redshift 1 is shown in Fig. 4.2, rescaled by (dt/dz) to estimate the maximal step width in Δz. The maximal eigenvalue for a proton only system is $\lambda_{\max} \approx 2 \cdot 10^3$, while for iron it is $\lambda_{\max} \approx 4 \cdot 10^5$. Note that these maximal eigenvalues roughly correspond to the maximal interaction rate. For a proton only system, the eigenvalues imply a stepsize of $\Delta z \lesssim 10^{-3}$. This can be done within reasonable computation time. However, in the case of heavier nuclei the eigenvalues are significantly larger, requiring a stepsize of $\Delta z \approx 10^{-6}$. For an integration from $z = 1$ to $z = 0$ this requires 10^{-6} integration steps. Equation systems with such a wide range in eigenvalues are called stiff. Even assuming only the computation time for the product $\Phi \times \vec{Y}$ of 2 ms and ignoring the re-computation of the elements in Φ, this implies a total integration time of more than 1000s for an iron system. Reducing the stepsize in an explicit integration scheme is therefore not an viable option for the propagation of heavy nuclei. One solution for this problem is to reduce the stiffness of the system. This means to integrate out the fast eigenvalues, i.e. the interactions of the respective particle species would be treated as instant with the secondaries directly injected into the system. This approach is used in MCEQ code [226] for atmospheric cascades, which only have a single injection of species at the start of the cascade and fixed interaction rates. However, in UHECR propagation both the interaction rates and the injection are time dependent. Specifically, the injection from local sources is of interest, as particles from these can reach the observer even for high interaction rates. Integrating out the eigenvalues for this system therefore requires a lot of manual fine tuning. PRINCE uses an implicit integration scheme instead.

---

[2] Note that the matrix inversion of Φ in 4.8 is computationally expensive. Therefore even for a system that reaches the steady state, numerical integration will be more effective in most cases.

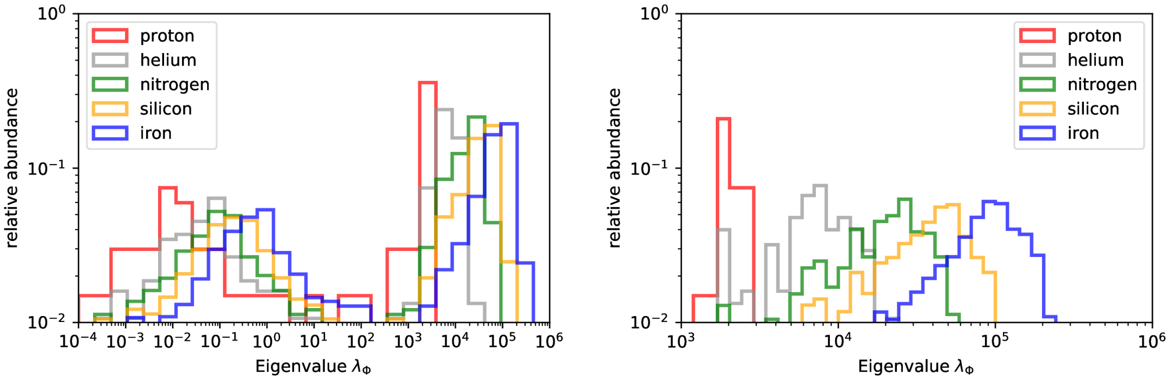

Figure 4.2: Distribution of eigenvalues of the coupling matrix $(dt/dz)\Phi$ for different injection species at $z = 1$. The left panel shows the full range of eigenvalues, while the right panels shows only the large eigenvalues. The counts are normed to 1 for each species. The eigenvalues cover several orders of magnitude, signaling that the problem is stiff. The maximal step size is determined by the largest eigenvalues. These are: proton: $2 \cdot 10^3$, helium: $1.7 \cdot 10^4$, nitrogen: $5 \cdot 10^4$, silicon: $1.3 \cdot 10^5$, iron: $3.4 \cdot 10^5$.

Implicit integration schemes are stable even for larger step sizes. The simplest one is the implicit Euler method, for which an integration step is given by:

$$\vec{Y}_{n+1} = \vec{Y}_n + \Phi(t_{n+1}) \times \vec{Y}_{n+1}\Delta t + \vec{J}(t_{n+1})\Delta t \tag{4.13}$$

$$= (1 - \Phi\Delta t)^{-1} \times (\vec{Y}_n + \vec{J}\Delta t) \tag{4.14}$$

this scheme does not lead to instabilities, as $(1 - \Phi\Delta t)^{-1}$ is always bounded, if the eigenvalues of Φ are negative. For this reason implicit integration schemes are suited for integration of stiff problems.

However, solving Eq. (4.13) formally requires an inversion of the interaction matrix in every step, which is computationally expensive. This can be overcome by approximate inversion techniques. Different implicit integration schemes are implemented in the SCIPY.INTEGRATE package. We choose the backwards differentiation solver (SCIPY.INTEGRATE.BDF). It is a more sophisticated version of the scheme in Eq. (4.13), choosing step width dynamically and using more previous timesteps (at t_{n-1}, t_{n-2}, ...) to better approximate the derivative at t_n. The inversion of Φ is done effectively by solving it iteratively using a variant of the Newton–Raphson method. Using these methods, allows for timesteps of $\Delta z \approx 10^{-3}$ for iron propagation.

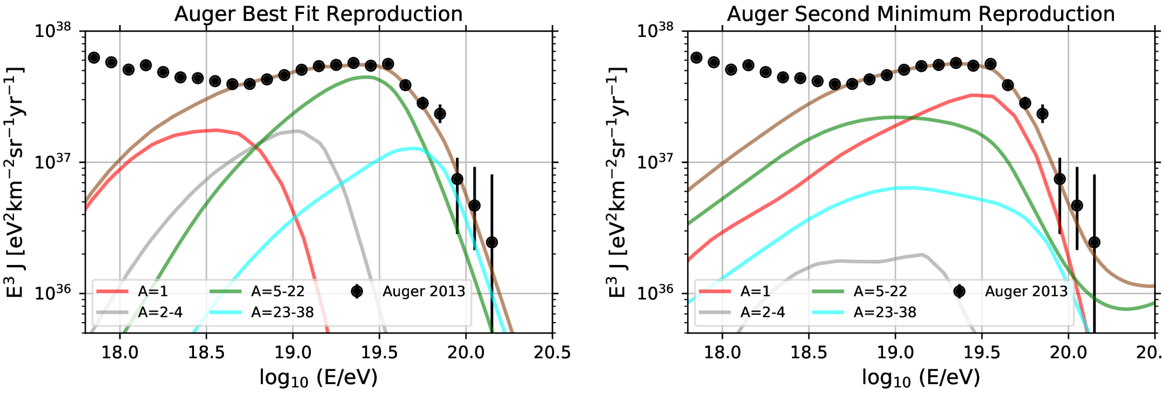

Figure 4.3: Cross-check for PriNCe reproducing the best fit and second minimum from the Auger combined fit [148] for the exact same input parameters. Compare to their Fig. 3 and Fig. 4, the color scheme was deliberately chosen to match. The original results were produced using SimProp. Using our method they are reproduced well.

4.5 Current implementation and selected results

The methods explained above are implemented in Prince in an object oriented way. The code currently contains the disintegration models PSB, Talys and Peanut. At higher energies, cross sections sampled from Sophia are used for photo-meson production. Other models can be included by supplying tabulated cross section data for the different particle channels. The models for different energy ranges are automatically combined by the code. We also include several EBL models, which can either be tabulated as a function of redshift or use an analytic redshift scaling.

PriNCe is able compute the propagation from $z = 1$ for one spectrum of UHECRs with composition including iron in $20 - 40$s depending on the number of stable particles contained in the disintegration model. It has no effect on performance whether the injection contains a pure or mixed composition. However, the computation time is greatly reduced if the maximal injected mass is lighter, as the system then contains less particle species. For protons only the propagation can be computed in a few hundred milliseconds.

To test the accuracy of the code, it was compared to different results from the literature [148, 210, 227] as well as directly to computations using SimProp. The result are reproduced quite well. Minor differences are expected due to the different methods. There are some smaller deviations even between the two established codes CRPropa and SimProp [153, 210]. As an example we show the reproduction of the Auger combined fit [148] in Fig. 4.3 and of Alves Batista et al. [210] in Fig. 4.4. The reproduction in Fig. 4.3 was done with Prince using the exact same

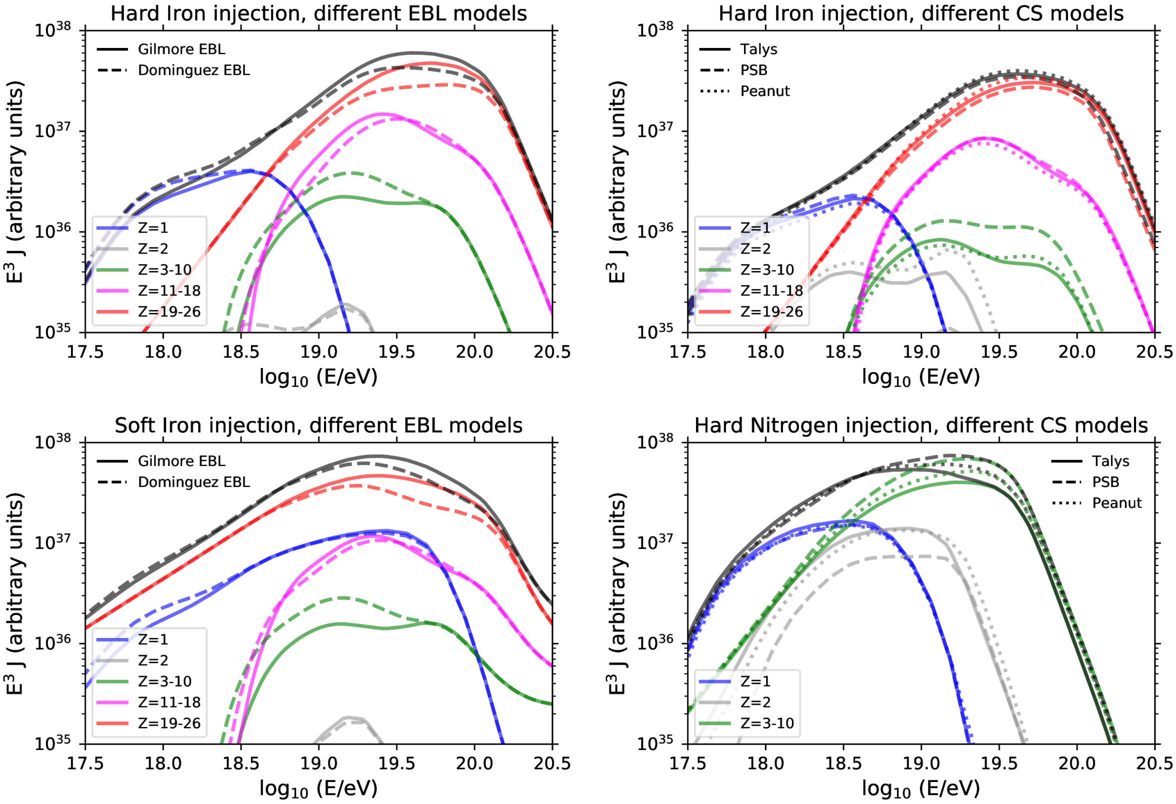

Figure 4.4: Cross-check for PriNCe reproducing propagation results for different cross section and EBL models Alves Batista et al. [210] using the exact same input parameters. Compare to their Fig. 6, Fig. 7 and Fig. 8, the color scheme was deliberately chosen to match. Both the spectral shapes and the model dependence are reproduced well.

injection parameters as reported for the best fit and second minimum in the Auger combined fit. The reproduction of Alves Batista et al. [210] in Fig. 4.4 illustrates the effect of different EBL and cross section models on the propagated spectrum for a pure composition of the injection spectrum.

4.6 Future optimizations and extensions

While the computation speed of PriNCe is already competitive compared to other codes, there are a few ways to further improve it. As mentioned before, we currently use the scipy package for the sparse matrix algebra and the numerical integrator. Implementing these methods on a lower level for our problem, the performance can therefore still be improved. We also do not yet include several physical processes, which should be implemented in the future.

Performance optimizations For standard linear algebra, so-called BLAS-libraries (*Basic linear algebra subprogram*) implement the most common operations on a very low computational level. A popular example is the Intel-MKL (*math kernel library*) [228]. Using these leads to a significant boost in performance even compared to low level implementations in languages like C and Fortran. While the dense matrix operations in many popular numpy-distributions are linked to Intel-MKL by default, the sparse matrix algebra in scipy relies on a C implementation. Intel-MKL, however, contains packages for sparse matrix algebra. By linking the most performance intensive operations manually, we could achieve a large performance boost with relatively little effort. Additionally, there are libraries optimized for graphics cards (e.g. cuBLAS and cuSparse[229]). Using these can further boost performance.

The code currently uses the standard scipy.BDF solver. Due to the generic interface which is intended for a wide range of problems, this solver does some redundant steps. For example, it internally re-approximates the Jacobian of Eq. (4.13) and tries to adaptively change the integration step size. This can be optimized by reimplementing a backwards differentiation solver specifically for our problem. This will also be necessary to have the internal integration steps calculated by the sparse matrix methods of MKL. The computation time could be further boosted by implementing the solver in a way that allows to propagate more than a single spectrum at a time. This can effectively boost the performance per spectrum, as the re-calculation of interaction only needs to be done once per redshift step.

Additional physics processes At the moment, PriNCe allows for the extragalactic propagation of UHECR nuclei accounting the interaction with photon backgrounds and the production of secondary neutrinos, photons and electrons. However, it does not yet contain interactions for electrons and photons. It can be extended by including several more processes in the future.

We currently do not include the electromagnetic cascade for photons and electrons. There are several processes which would need to be implemented for this. These have been discussed in the context of AGN spectra calculations in Gao et al. [230]. They are implemented in the C++-code AM^3, developed originally for AGN. The code is currently being adapted to GRBs and coupled to the nuclear cascade in NeuCosmA by our group. Collaborating on this to implement the physics also in PriNCe should therefore be possible with reasonable effort. In terms of computational performance, this should have little impact as the coupling between photons and electrons would barely increase the density in the interaction matrix Φ. This can be seen in Fig. 4.1, where the couplings for photon, electron and positron interactions would show up in the gray-shaded triangle.

Another extension, in the context of UHECR propagation, would be to include the diffusion in extragalactic magnetic fields. Simulating three dimensional diffusion in our method would be very performance intensive and is probably more efficiently done in Monte-Carlo codes. However, magnetic diffusion can be effectively described by a one-dimensional diffusion term. See e.g. Refs. [231–233] for examples. To include such a diffusion term one has to introduce a second grid in space x with the density at each point $Y_i(E, x)$, which increases the degree of the equation system to $N_{\text{particle}} \times N_{\text{E-grid}} \times N_{\text{x-grid}}$. For the exact methods described here, this would greatly increase the computation time. However, diffusion does not couple different species, keeping the system sparse. Secondly, the interaction rates are not explicitly dependent on x. Both these properties can be used to optimize the computation.

We developed PriNCe primarily for propagation of UHECRs. However, simulations of source physics in an isotropic approximation are very similar. PriNCe would therefore only need relatively little extensions to be applied to isotropic sources. Additional requirements are particle escape (which is trivially implemented by an additional escape rate Γ_{esc}) and synchrotron losses (which can be implemented analogously to pair-production losses). For static photon fields, the computation time should be much faster than for UHECR-propagation. Even for time-dependent photon fields the computation time should be faster, as the shorter timescales in sources require fewer integration steps. The implementation of a source model is however mainly interesting in terms of applications, if we can include the full electro-magnetic cascade. A code that can treat both nuclear interactions and the electro-magnetic cascade in sources self-consistently and in an efficient way would be novel and open up many new applications.

Chapter 5

Impact of model assumptions on interpreting UHECRs

In this chapter, the newly developed propagation code PriNCe, described in Chapter 4, is applied for a comprehensive study of a fit to the UHECR spectrum and composition. For this fit, we assume a generic source class and derive its general properties required to fit the observed Auger spectrum and composition. In this context, we discuss in detail the systematic effects of the different disintegration and air-shower models that were described in Chapter 3. This is done for the most recent Auger dataset (2017) [88, 93]. A similar study was published by the Auger Collaboration [148]. This combined fit (from here CF) was done for a generic source model with flat evolution. The impact of model assumptions was only discussed for a subset of the parameter space. We deliberately use the same framework, but extend it to include source evolution as an additional free parameters. A reproduction of the main CF results is presented in the Appendix B.1 to discuss the impact of the updated (2017) data set to the one used in the CF (2015). From our results, we also derive the expected flux of cosmogenic neutrino and its robustness under model assumptions.

This chapter is based on the previous publication in Heinze et al. [3]. This study is the first application of the new PriNCe propagation code. I also implemented the tools for the computation of the parameter space and the fit to data. With these, I performed the parameter scans and evaluated them in discussion with the other authors.

5.1 Generic source model assumptions

As discussed in Section 3.1, the expected spectrum from Fermi-acceleration is a power-law ($\propto E^{-\gamma}$) with spectral index of $\gamma \approx 2$. However, this is only the spectrum at the acceleration site. If charged particles are magnetically confined, the magnetic fields act as a high-pass filter, making the escaping spectra significantly harder. Therefore, spectra of escaping charged particles that significantly differ from E^{-2} are not unexpected for a single source.

Models of UHECR propagation therefore usually assume a simple parameterization of the spectrum ejected from the source, where the spectral index and a maximal energy cutoff are free parameters. This parameterization then represents the total ejection from a generic population of sources. The exact shape of the cutoff is again dependent on the source environment. A common choice is to parameterize it by an exponential cutoff. However, this choice is in principle arbitrary, since the exact form would depend on the details of acceleration and escape. We choose the exact same parameterization as in the CF [148] (see also e.g. Romero-Wolf and Ave [151], Alves Batista et al. [153]) to make the results directly comparable:

$$J_A(E) = \mathcal{J}_A \; f_{\mathrm{cut}}(E, Z_A, R_{\mathrm{max}}) \; n_{\mathrm{evol}}(z) \left(\frac{E}{10^9 \;\mathrm{GeV}}\right)^{-\gamma} , \tag{5.1}$$

for each the nuclear species A, with the five injection elements: ^{1}H, ^{4}He, ^{14}N, ^{28}Si and ^{56}Fe, which cover roughly equal distances in $\ln(A)$. The different elements share a common spectral index γ and maximal rigidity $R_{\mathrm{max}} = E_{\mathrm{max}}/Z_{\mathrm{A}}$. $\mathcal{J}_A$ are the normalizations of the spectra of each of the injection elements, representing the number of particles ejected from the sources per unit of time, comoving volume and energy. These fractions are free parameters and will be determined by the fit to spectrum and composition. For the cutoff, we again adopt the definition of the CF:

$$f_{\mathrm{cut}}(E) = \begin{cases} 1 & , E < Z_{\mathrm{A}} R_{\mathrm{max}} \\ \exp\left(1 - \dfrac{E}{Z_{\mathrm{A}} R_{\mathrm{max}}}\right) & , E > Z_{\mathrm{A}} R_{\mathrm{max}} . \end{cases} \tag{5.2}$$

The definition in Eq. (5.2) is convenient for defining mass fractions, as the spectra are parallel power-laws at energies below R_{max}. Still, the definition of the injected mass fractions is ambiguous. In the CF, the injected fractions f_{A} are defined at a fixed energy point (10^9 GeV), relative to the total spectrum. This definition is straightforward to obtain from $\mathcal{J}_A$ as $f_A = \mathcal{J}_A / \sum_A \mathcal{J}_A$. This definition was chosen for the CF as it is easily comparable between different studies. However it has a strong degeneracy with γ and R_{max}, i.e. for soft spectra the fraction of protons at low energy can be very high without contributing much at the relevant energies above the *ankle*. While the iron fraction typically seems very low in f_A, it can actually be significant due to its higher maximal energy (for equal R_{max}). f_{A} therefore gives a somewhat misleading impression of the actual contribution of elements. We chose to define the mass fractions by the fraction of the integrated energy density, as this is physically more meaningful:

$$I_A^9 = \frac{\int_{E_{\mathrm{min}}}^{\infty} J_A(E) E \mathrm{d}E}{\sum_A \int_{E_{\mathrm{min}}}^{\infty} J_A(E) E \mathrm{d}E} , \tag{5.3}$$

where we choose $E_{\mathrm{min}} = 10^9$ GeV as the lower boundary. We will mostly refer to I_A^9 in our study, while still providing the f_A for comparability with the CF.

Eq. (5.1) contains a function $n_{\text{evol}}(z)$ to account for the source evolution with redshift. In the CF only flat evolution, $n_{\text{evol}}(z) = 1$, was considered. We choose the simple parameterization:

$$n_{\text{evol}}(z) = (1 + z)^m. \tag{5.4}$$

For the variable m, the function approximates all known continuous source density functions within the UHECR horizon $z \lesssim 1$. However the production of secondary messengers, like neutrinos, at higher redshift is affected by the source evolution. We will discuss the impact of different extrapolations of the source evolution at higher redshift, when estimating the cosmogenic neutrino flux in Section 5.4,

We inject five representative injection elements: hydrogen (^{1}H), helium (^{4}He), nitrogen (^{14}N), silicon (^{28}Si) and iron (^{56}Fe). These mass groups cover the space in $\ln(A)$ roughly equally, which is the main quantity scaling the observed X_{max}. We verified that choosing different injection elements in the same mass groups yields qualitatively similar results. An exception is carbon (^{12}C), which disintegrates very efficiently into ^{4}He, compared to ^{14}N or ^{16}O, reducing the fraction of ^{4}He required directly from the sources. This is only true for the disintegration models PEANUT and TALYS, as PSB does not contain ^{4}He emission. We chose to use ^{14}N to represent this mass group, which is the same as in the CF.

This generic parameterization covers the most important features of the spectra expected from different UHECR-source candidates. However, a few relatively strong assumptions about the source candidates are still implied. The assumption of rigidity-dependent maximal energy applies only to sources with lower luminosity, as discussed in Section 3.1. For high luminosities, the maximal energy is typically limited by interactions, which results in similar maximal energies for the different nuclear species [118, 124]. We will discuss the impact of the rigidity dependent maximal energy in Section 5.3.3.

The assumptions of an equal spectral index for all nuclear species is similarly affected by the luminosity. For higher luminosities the interactions produce neutrons, which are not magnetically confined and escape. Their subsequent escape produces a proton component with softer spectral index than the spectra of directly escaping charged particles. Additionally, very hard spectral indices ($\gamma < 0$) might be difficult to reconcile with the assumption of a population of a large number of sources. While they can be expected from a single source depending on the acceleration and escape mechanisms, all sources of a population would have to accelerate to the same maximal energy. Otherwise the superposition of the individual spectra leads to a smoothed-out total spectrum [139, app. E]. Still, we restrict ourselves to the generic propagation model in Eq. (5.1) for the main results, as it is flexible and covers many features of a wide range of different source models.

5.2 Simulation and fitting procedure

In this section, we summarize the relevant setup for our simulations and parameter scans. The generic source model has eight free parameters: $R_{\max}$, γ, m and free normalizations $\mathcal{J}_A$ corresponding to the five injection elements. We allow for a shift $\delta_{\rm E}$ in energy within the systematic uncertainty given by Auger (±14%) [88]. The UHECR transport equation (Eq. (3.3)) is linear in the normalization factor $\mathcal{J}_A$ but not in the other source parameters (γ, $R_{\max}$ and m). This means that every change of γ, $R_{\max}$ or m needs a full recomputation, while for $\mathcal{J}_A$ the result can simply be rescaled. We therefore use a discrete grid in γ, $R_{\max}$ and m with the following ranges and stepsizes:

	min	max	stepsize
γ	-1.5	2.5	0.05
$\log_{10}(R_{\max})$	9.7	11.7	0.05
m	-6	6	0.2

For each point of this three-dimensional (3D) source-parameter grid, we separately compute the spectra at Earth for the five injection elements ($\sim 1.5 \cdot 10^6$ individual simulations for one choice of the photo-nuclear interaction model). Since the propagated spectra are linearly dependent on the injection norms $\mathcal{J}_A$, the all-particle spectrum can be calculated as a linear superpostion of the propagated spectra obtained for single element injection.

For each point on the 3D grid, we fit the nuclear fractions $\mathcal{J}_A$ and energy shift $\delta_{\rm E}$ to the spectrum and the first two moments of $X_{\max}$ for each triplet in $(R_{\max}, \gamma, m)$ using the Minuit package [234][1]. The translation from individual mass spectra at the top of the atmosphere to $\langle X_{\max} \rangle$ and $\sigma(X_{\max})$ is performed with the parameterization from Abreu et al. [19] described in Section 3.3, using updated parameter sets[2] for Sibyll 2.3 and Epos-LHC.

We use the following χ^2 definition to estimate goodness of fit:

$$\chi^2_{\mathcal{F}} = \sum_i \frac{(\mathcal{F}(E_i) - \mathcal{F}_{\rm model}(E_i, \delta_E))^2}{\sigma_i^2}, \tag{5.5}$$

where $\chi^2_{\mathcal{F}}$ refers to each of the three observables $\mathcal{F}$, namely the combined energy spectrum, $\langle X_{\max} \rangle$ and $\sigma(X_{\max})$ [88, 93]. We minimize the total χ^2, which is the sum of the fit to all three observables. A nuisance parameter δ_E captures the uncertainty in the energy scale. We assume its distribution to be flat within ±14%. The fit takes into account all data points above

---

[1] We use the iMinuit interface `https://github.com/iminuit/iminuit`.

[2] Private communication with S. Petrera

$E_{\mathrm{min}} = 6 \cdot 10^9$ GeV, i.e. all data points above the *ankle*. The global best fit χ^2_{min} is found by minimizing over all points of the 3D parameter grid.

The full grid in $\gamma - R_{\mathrm{max}} - m$ has to be recomputed for each disintegration-model, as they change the interaction rates during propagation. The air-shower model only converts the propagated spectra's $\ln(A)$-distribution to X_{max}, therefore only the fit to data has to be repeated for a changed air-shower model. In all cases, we use Gilmore et al. [194] as the EBL model. The impact of the uncertainty in the EBL model is small compared to the disintegration and air-shower model, as discussed in Alves Batista et al. [210]. To find the χ^2 values for the UHECR fits within the entire parameter space, the simulations are performed starting from redshift $z_{\mathrm{max}} = 1$. Once the 3σ confidence intervals are localized, we run additional simulations starting from $z = 3$ to compute cosmogenic neutrino fluxes. We verified that these simulations at higher redshift give the same results for the UHECR spectra, as expected due to the UHECR horizon at $z = 1$. This is demonstrated in Section 6.3.

We then use $\Delta\chi^2 = \chi^2 - \chi^2_{\mathrm{min}}$ to draw contours around the best fit point by projecting to planes of two parameters by minimizing over all other parameters of the scan. This frequentist approach is sufficient to draw contours and discuss the correlations among the continuous source parameters. However, it cannot be as easily applied to discrete physical parameters choices, such as the combinations of models for the photon background, the disintegration and the hadronic interaction model. For example, Romero-Wolf and Ave [151] parameterized the hadronic model as a linear superposition of Epos-LHC and Sibyll 2.3 with a continuous nuisance parameter. However such a nuisance parameter is neither physically meaningful nor unbiased. Instead it hides the qualitative differences resulting from these two different models. We therefore choose to show the different discrete model combinations and discuss their qualitative differences in the fit contours.

5.3 Results of the UHECR fit

To verify our methods, we used them to directly reproduce the results of the CF. Details on this reproduction are given in Appendix B.1. The main results are well reproduced, with some minor differences due to details of the propagation codes and fitting procedures. With the updated 2017 data set, a small but significant iron fraction is required to fit the spectrum. This is a direct effect of the increased statistics at the cutoff, as is discussed in Appendix B.1 by comparing to the older 2015 data set.

In the following, we will discuss our UHECR fit in context of different disintegration and air-shower models. However, we first discuss the main features for the “baseline” case chosen to be the combination of Talys as disintegration model and Sibyll 2.3 as air-shower model

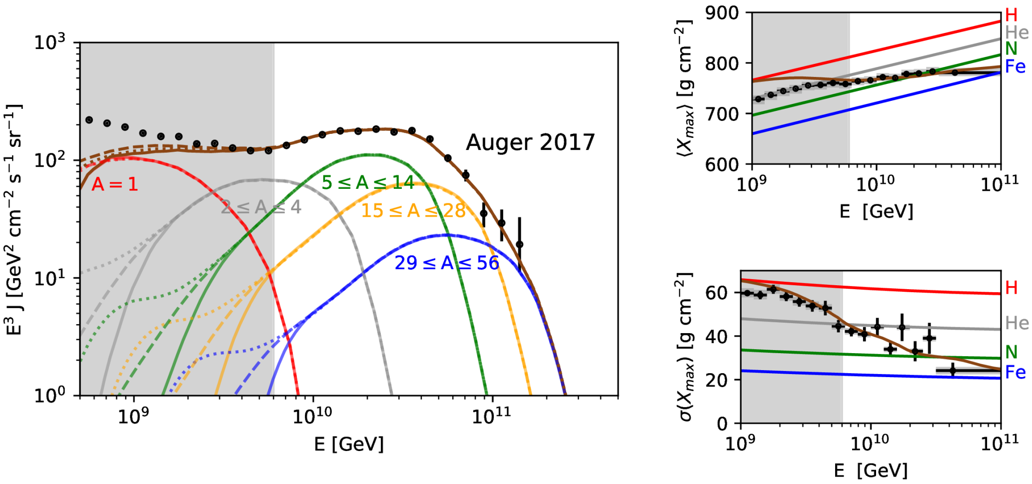

Figure 5.1: Spectrum (left panel) and composition observables (right panels) corresponding to the best fit to the Auger 2017 data, for the baseline model combination Talys and Sibyll 2.3. The corresponding injection at the source is found in Fig. 5.2. The dashed and dotted curves show the spectrum for propagation from $z_{\max} = 3$ with different continuations of the source evolution. The impact on the UHECR fit is negligible. Figure adapted from Heinze et al. [3].

in Section 5.3.1. It is then extended to the different model combinations in Section 5.3.2. In Section 5.3.3 we will discuss the impact of the rigidity dependence by using a generalised cutoff function. The impact of the model choices on the injected composition is discussed in Section 5.3.4.

5.3.1 Baseline model combination Talys – Sibyll 2.3

The combination of Talys as disintegration and Sibyll 2.3 as air-shower model serves as the 'baseline' model, simply because it gives the lowest χ^2 out of the more realistic disintegration models. We will used this baseline model to establish the main parameter correlations of the fit. The other model combinations are discussed in Section 5.3.2. Before discussing the correlations in the full parameter space, we show two distinct examples from the parameter space: The best fit ($\chi^2/\text{dof} = 27.0/21$) is found at strong source evolution and hard spectral index and a second example chosen with negative source evolution and softer spectral index, which also leads to reasonably good fit ($\chi^2/\text{dof} = 34.6/21$). The parameters of the two points are given in Tab. 5.1.

The spectrum and composition of the best fit in our baseline model are given in Fig. 5.1. The propagated spectra for different mass groups A are highlighted. Both spectrum and composition fit well. Below the fit-range, the propagated spectrum cannot fully sustain the flux and is

Table 5.1: Best fit parameters for the 3D parameter scan with free source evolution for the baseline case of the combination Talys - Sibyll 2.3. Uncertainties are given for the 1σ-interval (for 1 d.o.f.). The left column shows the result for a fixed point within the 3σ-contours at negative source evolution and soft spectral index. In Fig. 5.1, these two points are marked by ● and ⋆ respectively.
Table extended from Heinze et al. [3].

	Talys - Sibyll 2.3					
	best fit (●)			second example (⋆)		
γ	$-0.80^{+0.27}_{-0.23}$			1.5 (fixed)		
$R_{\max}$ (GV)	$(1.6 \pm 0.2) \cdot 10^9$			$4 \cdot 10^9$ (fixed)		
m	$4.2^{+0.4}_{-0.6}$			-5.0 (fixed)		
δ_E	$0.14^{+0.00}_{-0.03}$			$0.14^{+0.00}_{-0.01}$		
f_A(%)	H	He	N	H	He	N
	$0.0^{+42.6}_{-0.0}$	$82.0^{+3.8}_{-6.4}$	$17.3^{+1.0}_{-1.1}$	$0.0^{+3.4}_{-0.0}$	$26.5^{+6.1}_{-7.3}$	$61.9^{+2.3}_{-2.6}$
	Si	Fe		Si	Fe	
	0.6 ± 0.1	$(2.0 \pm 0.8) \cdot 10^{-2}$		$10.6^{+2.6}_{-2.8}$	$0.9^{+0.6}_{-0.7}$	
I_A^9(%)	H	He	N	H	He	N
	$0.0^{+1.2}_{-0.0}$	$9.8^{+2.8}_{-2.9}$	$69.2^{+1.5}_{-1.6}$	$0.0^{+2.5}_{-0.0}$	$24.1^{+5.8}_{-6.8}$	$62.4^{+2.3}_{-2.6}$
	Si	Fe		Si	Fe	
	$17.9^{+3.2}_{-3.5}$	$3.2^{+1.2}_{-1.3}$		$12.4^{+3.0}_{-3.2}$	1.0 ± 0.7	
χ^2 / dof	27.0 / 21			34.6 / 21		

dominated by protons, which are too light compared to the observed composition. This lower energy component is expected to be suppressed by diffusion in extragalactic magnetic fields, which we did not include for our simulations. The generic source model therefore cannot account for the component below the ankle. For a fit with an extended range, one would have to include an additional (galactic) component. The dashed and dotted lines show the result for the same injection spectrum but with the source evolution continued to $z_{\max} = 3$. The effect on the fit is obviously minor, as expected due to the UHECR horizon at $z = 1$. This only has an effect on the cosmogenic neutrino flux, as will be discussed in Section 5.4.

The corresponding injection spectrum is shown in the left panel of Fig. 5.2. The right panel shows the propagated spectrum again, but this time with the secondaries grouped by the initial injected mass. The dashed lines show the spectra of the initial injection elements that survive the propagation without disintegrating. This clearly shows a pile-up effect from higher redshift. The injected spectrum is very hard ($\gamma = -0.8$), which clearly separates the injected mass spectra. The secondaries from disintegration pile up at lower energies, leading to softer total spectra with a stronger overlap of different masses. The best fit for the proton component is 0. The protons in the propagated spectrum are therefore only from disintegration. The upper limit on

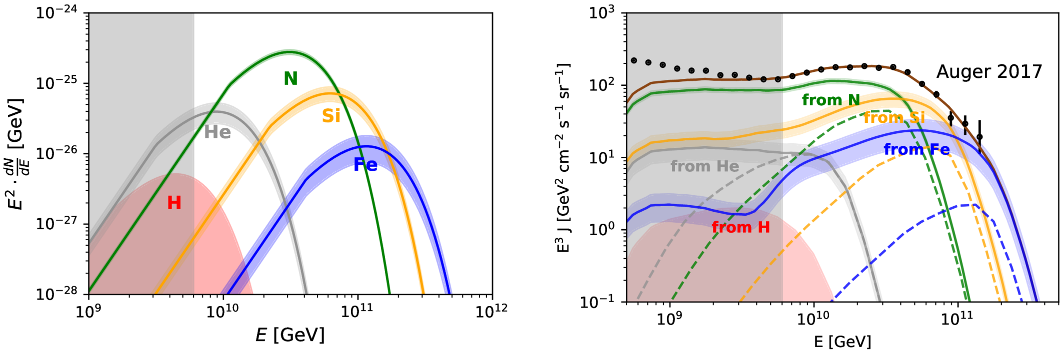

Figure 5.2: *Left:* Injection spectra for the five injected elements corresponding to the best fit for the 3D parameter scan in Fig. 5.4 ($\gamma = -0.8$, $R_{\mathrm{max}} = 1.6 \cdot 10^9$ GV, $m = 4.2$). The shaded regions indicate the 1σ uncertainties to the normalization of each injection corresponding to the fit (for $\gamma, R_{\mathrm{max}}, m$ fixed). While the best fit proton fraction is 0, there can be a significant proton contribution within the uncertainty. *Right:* The spectrum after propagation with secondaries grouped by the initial injection elements. The dashed curve shows the surviving spectrum of the respective element. Left panel taken from Heinze et al. [3].

the proton component (shaded region) only allows for a small component of directly injected protons.

Fig. 5.2 shows the same plot for the second example with negative source evolution and softer spectral index. The injection has a softer spectral index ($\gamma = 1.5$), which is closer to the index of the observed spectrum. Disintegration is much less efficient due to the more local sources. Therefore, there is a significant component of surviving injection elements. Consequently, there is less pile up of protons below the fit-range. Overall, the observed spectrum is more directly impacted by the shape of the injection function. Even though the number of secondary protons from disintegration is lower in this case, the component of directly injected protons is limited to a few percent.

The full parameter space is shown in Fig. 5.4, with the two discussed examples marked by $\bullet$ (best fit) and $\star$. The $\gamma - m$ parameter plane exhibits a clear anti-correlation, as already noticed for example in Refs. [147, 162]. As discussed at the two examples, positive source evolution ($m > 0$) result in a pile up from more distant sources, effectively softening the spectrum at Earth. This pile up is compensated by harder spectra at the source. Contrariwise, a high density of local sources ($m < 0$) allows for spectral indices almost compatible with Fermi acceleration. The result for this model combination favors positive source evolutions, covering star-forming objects, GRBs and Blazars. The 3σ contours still leave room for negatively evolving sources such as TDEs [130].

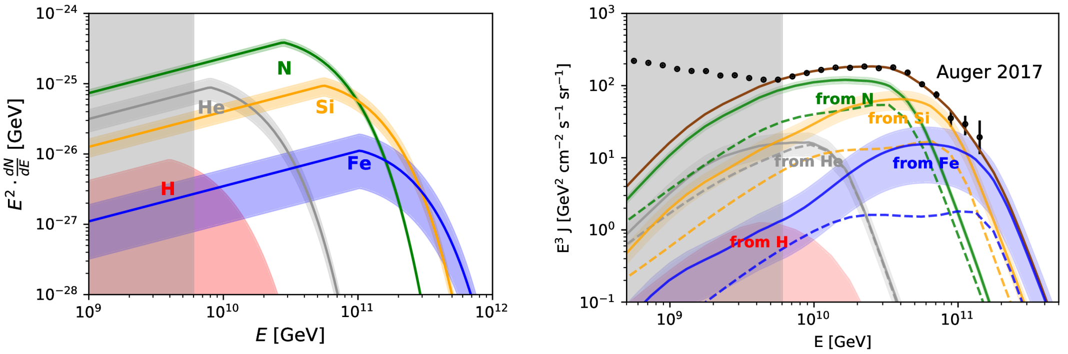

Figure 5.3: Same as Fig. 5.2, but for the second point at ($\gamma = 1.5$, $R_{\max} = 4. \cdot 10^9$ GV, $m = -5$), which is marked by $\star$ in Fig. 5.4.

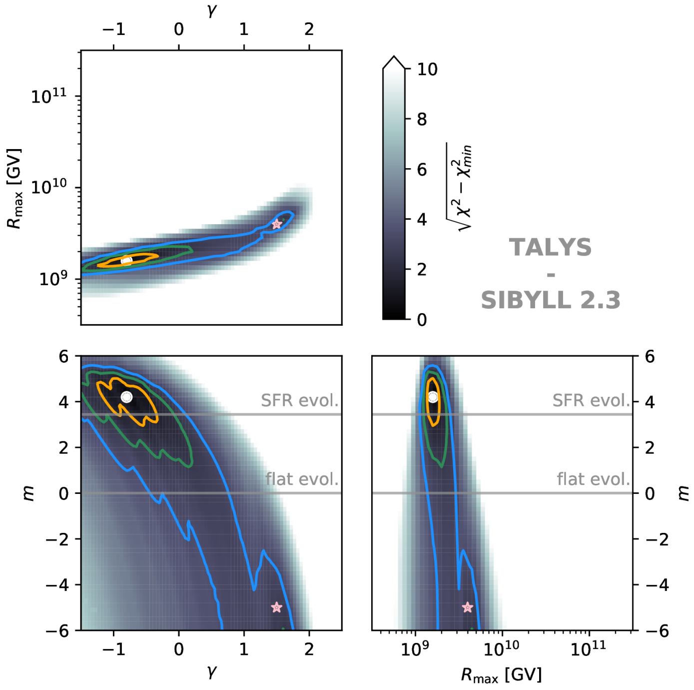

Figure 5.4: Parameter space in γ, $R_{\max}$ and m for the model combination Talys and Sibyll 2.3 (baseline case) with free energy scale as a nuisance parameter. The best fit, found at $\gamma = -0.8$, $R_{\max} = 1.6 \cdot 10^9$ GV and $m = 4.2$, is marked by a white dot. A second point at $\gamma = 1.5$, $R_{\max} = 4 \cdot 10^9$ GV and $m = -5$ is marked by a pink star. The colored shading corresponds to $\sqrt{\chi^2 - \chi^2_{\min}}$, while $\chi^2 - \chi^2_{\min}$ is used to determine the contours, which are given for $1\sigma, 2\sigma, 3\sigma$ (for 2 d.o.f.). In each 2D panel, the third parameter is treated as a nuisance parameter and minimized over to project the 3D parameter space. Figure taken from Heinze et al. [3].

The contour in the $\gamma - R_{\max}$ plane is similar to the flat evolution case in the CF. Although the $\gamma \approx 1$ corresponding to Fermi acceleration with diffusive escape is within the 95% contour, the preferred spectral indices result in flat or almost monochromatic spectra with $\gamma < 1$. The slight correlation in this plane comes from the definition of the spectrum above the cutoff, in which hard spectral indices can slightly compensate a lower $R_{\max}$.

The $R_{\max} - m$ plane also exhibits a low rigidity cutoff for every choice of the source evolution within the 95% CL. For low $R_{\max}$ there is less disintegration during propagation leading to smaller superposition of the individual mass spectra. This is required by the observed $\sigma\left(X_{\max}\right)$, as was discussed in Section 3.3, specifically Fig. 3.6. The low rigidity scenario means that the cutoff in the spectrum is due to the sources running out of energy, rather than due to photo-disintegration losses. This is expected for negative source evolution, where disintegration is less efficient and cannot fully account for the observed cutoff. Still, it also holds for strong positive evolution.

5.3.2 Impact of the different model combinations

With the basic parameter correlations established for the baseline case, we now discuss the differences due to the different choices of disintegration and air-shower models. The fit is repeated for all combinations of the disintegration models PSB, Talys and Peanut and the air-shower models Epos-LHC, Sibyll 2.3 and QGSjetII-04. The most relevant correlations are visible in the $\gamma - m$ plane. The different model combinations are compared in this plane in Fig. 5.5. The comparisons for the other two planes can be found in the Appendix in Fig. B.3 and Fig. B.4. The corresponding best fit parameters are also reported in the Appendix in Tab. B.2.

Consistent with what was found in the CF, we cannot find reasonable fits for QGSjetII-04. As discussed in Section 3.3, this is due to the model's broad $X_{\max}$ distributions, in combination with a small $\langle X_{\max} \rangle$, which is inconsistent with the observed $X_{\max}$ distribution. For all the other model combinations we find satisfactory best fits with $\chi^2/\mathrm{dof} \approx 1.4 - 2.0$. The shower model clearly has a stronger impact on the fit contours than the disintegration model, as can be seen comparing the columns in Fig. 5.5. The differences between Talys and Peanut are mainly quantitive. Compared to the other two models PSB prefers slightly stronger source evolution. Interestingly, for the PSB model in combination with Sibyll 2.3, negative source evolutions are excluded at 3σ. Note that there are significant differences in the injected mass fraction, which are minimized over in this figure. These are discussed in Section 5.3.4.

The anti-correlation between m and γ is found for all combination of disintegration and shower model (excluding QGSjetII-04). However, when exchanging Sibyll 2.3 with Epos-LHC, the 3σ contour in Fig. 5.5 is shifted towards more local sources and/or more monochromatic

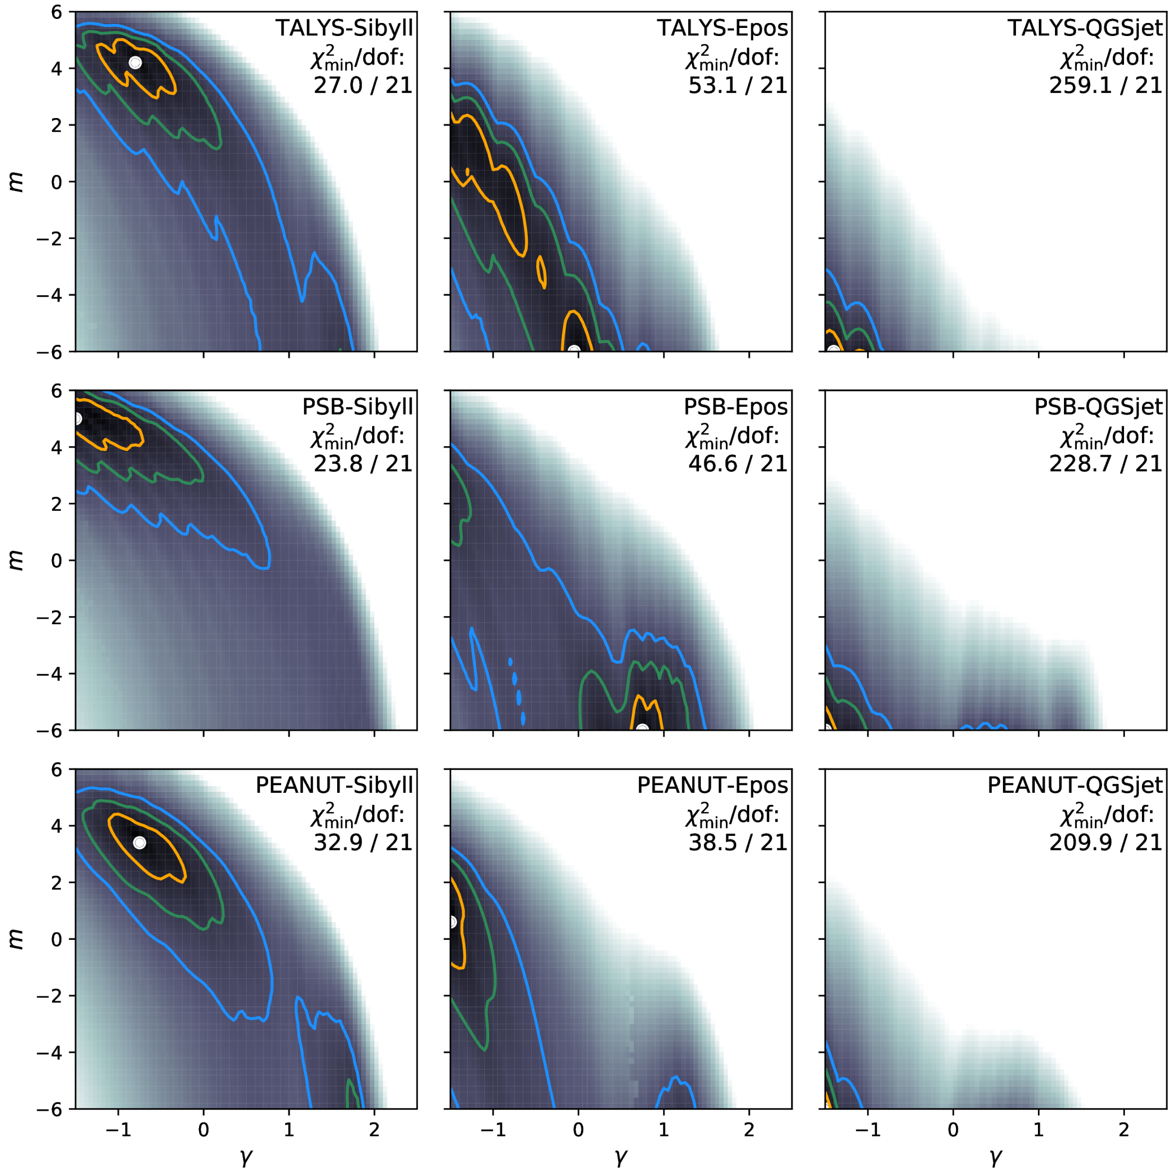

Figure 5.5: Parameter space in γ and m minimized over the third parameter $R_{\max}$ for different combinations of disintegration and air-shower models. The color code and contours are defined as in Fig. 5.4. Rows from top to bottom: Talys, PSB, Peanut. Columns from left to right: Sibyll 2.3, Epos-LHC, QGSjetII-04. The corresponding best fit parameters can be found in Tab. B.2 (appendix). Figure taken from Heinze et al. [3].

spectra. The reason for this is that Epos-LHC, compared to Sibyll 2.3, predicts less shower-to-shower fluctuation decreasing the $\sigma(X_{\mathrm{xmax}})$, while at the same time its $\langle X_{\mathrm{max}}\rangle$ predicts a lighter composition of the measurements. In combination, this allows for less overlap of individual mass spectra. This is satisfied by either harder spectra or reduced disintegration due to more local sources, as explained in the previous section. At the same time, the maximal rigidity R_{max} is more constrained for Epos-LHC than for Sibyll 2.3 (see Fig. B.4). This is again since the allowed overlap of mass spectra and therefore the rate of disintegration is more constrained for Epos-LHC.

The $\chi^2_{\mathrm{min}}/\mathrm{dof}$ is slightly worse when using Epos-LHC (≈ 2.0) compared to Sibyll 2.3 (≈ 1.4), mainly because the fit to $\langle X_{\mathrm{max}}\rangle$ is worse. The difference is however not strong enough to discriminate between these models, due to the large systematic uncertainties in X_{max} on the level of $\pm 10\,\mathrm{g\ cm^{-2}}$. We found that the differences in the fit for the two air shower models can be somewhat alleviated by shifting the X_{max} data within its systematic uncertainties. There are however strong bin-to-bin correlations for the X_{max} data, that cannot be easily treated without detailed information only available to the Auger collaboration. We therefore chose not to include the systematic X_{max} uncertainties in our fit.

5.3.3 Impact of the rigidity-dependent maximal energy

For the fits in the previous sections, we assumed that the maximal rigidity is the same for each injected element, i.e. $E_{\mathrm{max}} = Z_A \cdot R_{\mathrm{max}}$. This assumption is valid for optically thin sources for which the acceleration rate is limited by the size of the region, as discussed in Section 3.1. For higher radiation densities, however, the maximal energy of each element is controlled by its disintegration and energy loss rates. For high luminosity sources, the maximal energies for different elements can then be in a similar range with the details depending on the source environment [124]. To test the assumption of rigidity dependence, we introduce a fourth parameter (α) in the cutoff function Eq. (5.2), that controls the scaling of the maximal energy:

$$f_{\mathrm{cut}}(E) = \begin{cases} 1 & , E < {Z_{\mathrm{A}}}^{\alpha} R_{\mathrm{max}} \\ \exp\left(1 - \dfrac{E}{{Z_{\mathrm{A}}}^{\alpha} R_{\mathrm{max}}}\right) & , E > {Z_{\mathrm{A}}}^{\alpha} R_{\mathrm{max}} \end{cases} . \tag{5.6}$$

For $\alpha = 1$ this definition retains the definition of the result in the previous three-dimensional fits. For $\alpha = 0$ all elements have the same maximal energy. Realistic source models will likely be in the range $0 \leq \alpha \leq 1$. Still, we extend the range for the parameter scan to $-0.5 \leq \alpha \leq 1.5$ to get a full coverage of the allowed parameter space. With this additional parameter we now repeat the fit following the same procedure as before. To save computation time, we decrease the

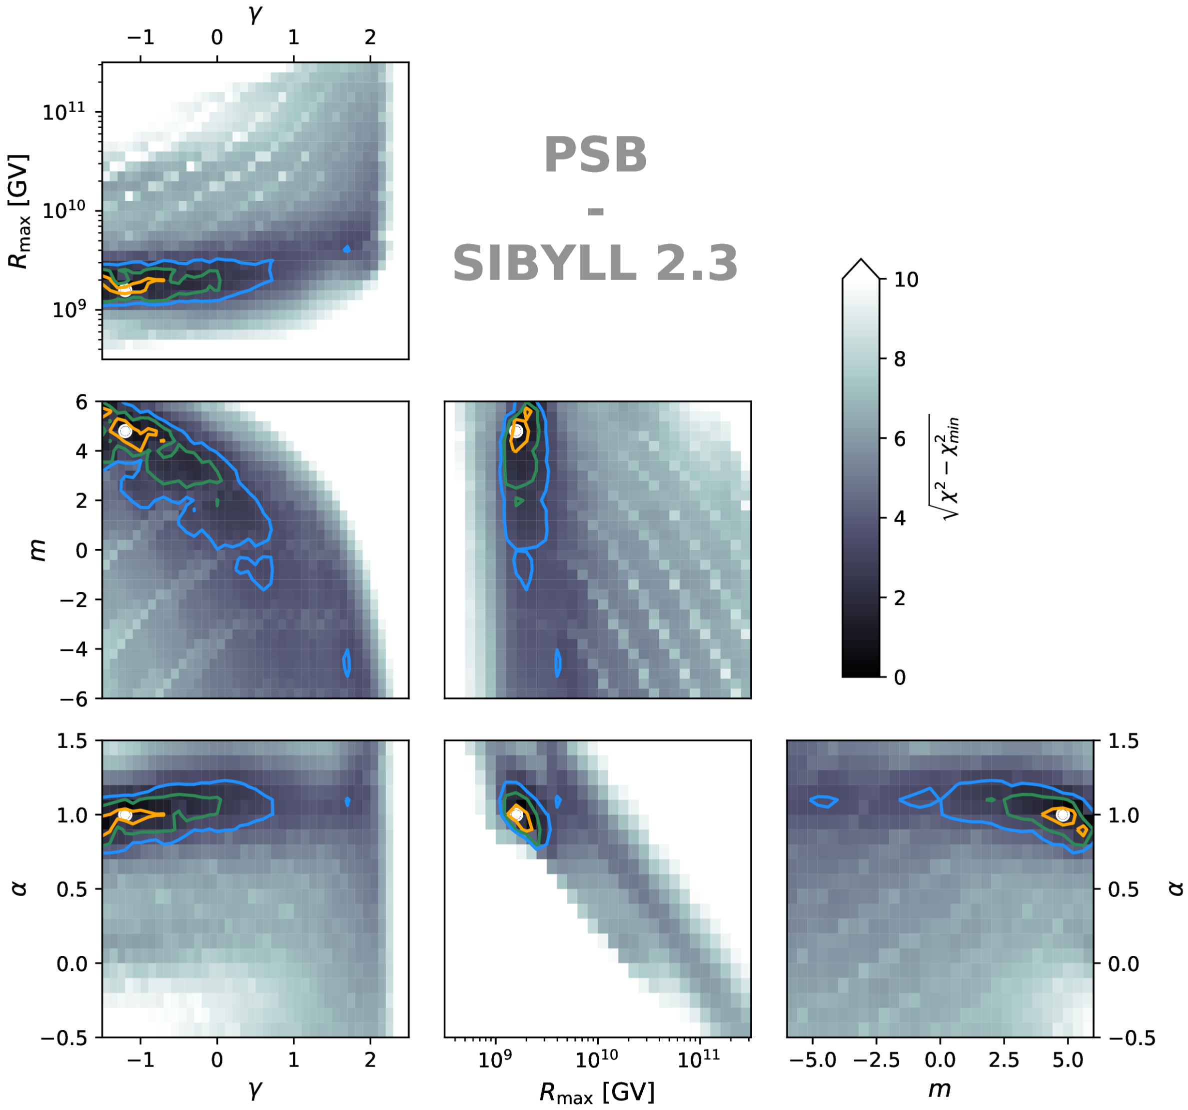

Figure 5.6: Parameter space in γ, R_{max}, m and α for the model combination PSB and Sibyll 2.3. This is the same as Fig. 5.4, but for the PSB disintegration model and extended by a fourth parameter α that controls the scaling of the maximal energy according to Eq. 5.6. The best fit at $\gamma = -1.2^{+0.2}_{-0.3}$, $R_{\mathrm{max}} = 1.6^{+0.4}_{-0.0} \cdot 10^9$ GV, $m = 4.8^{+0.8}_{-0.4}$ and $\alpha = 1.0^{+0.0}_{-0.1}$ is marked by a white dot.

resolution in each parameter, computing only $41 \times 31 \times 31 \times 21$ points in $\gamma \times R_{\mathrm{max}} \times m \times \alpha$. We also choose only the model combination PSB - Sibyll 2.3, since PSB takes less computation time for propagation due to its fewer disintegration channels.

The result is shown in Fig. 5.6. The best fit is found at $\gamma = -1.2^{+0.2}_{-0.3}$, $R_{\mathrm{max}} = 1.6^{+0.4}_{-0.0} \cdot 10^9$ GV, $m = 4.8^{+0.8}_{-0.4}$ and $\alpha = 1.0^{+0.0}_{-0.1}$. These ranges are fully consistent with the best fit found in the previous three-dimensional fit (compare to Appendix, Tab. B.2). The allowed range in α is very narrow around 1. The other contours therefore retain roughly the shapes of the previous fit, where $\alpha = 1$ was fixed. This means that the choice of rigidity-dependent maximal energy does not imply a strong restriction on the parameter space. There are some correlations visible outside of the 3σ contours. There is, for example, an anti-correlation in $R_{\mathrm{max}} - \alpha$, which is expected since $E_{\mathrm{max}} = Z^{\alpha} \cdot R_{\mathrm{max}}$ needs to fit the cutoff in the observed spectrum. However, these regions do not give a good fit to the observed data anyway.

The data strongly favors a rigidity-dependent cutoff due to the small overlap of masses allowed by the combination $\langle X_{\mathrm{max}} \rangle$ and $\sigma(X_{\mathrm{max}})$, as discussed in Section 3.3. Epos-LHC allows for even less superposition of masses compared to Sibyll 2.3. In fact, for the combination of PSB - Epos-LHC we find the best fit at $\alpha = 1.2$, which further reduces the overlap of masses. However, values of $\alpha > 1$ are not consistent with the standard acceleration mechanisms. We did not test the impact of other disintegration models than PSB explicitly, since this would have required additional computation time. However, similar results are expected for the other two models, since more detailed disintegration chains will only increase the overlap of masses and therefore also require small initial overlap. With the current data and hadronic models the assumption of rigidity dependence is a good ad-hoc choice for generic models. In any case the parameterization with α chosen here is rather simple. For a more physically meaningful definition of the scaling of E_{max}, one could define it by the disintegration rate on a generic photon spectrum. However, this would require at least some assumption about the sources of UHECRs and lead to a definition that is not easily described by a single parameter.

5.3.4 Injected composition

For the parameter scans presented in the previous sections, we minimized over the injected mass composition. The injected mass composition permitted by Auger data is however relevant when developing source models. While the composition at observation is fixed (within the uncertainty of air-shower models and data), it can have significantly different interpretations in terms of the composition ejected from the source. For example, more distant sources have to inject a heavier composition, as disintegration is more efficient during propagation. The inferred composition therefore depends both on the disintegration and the air-shower model.

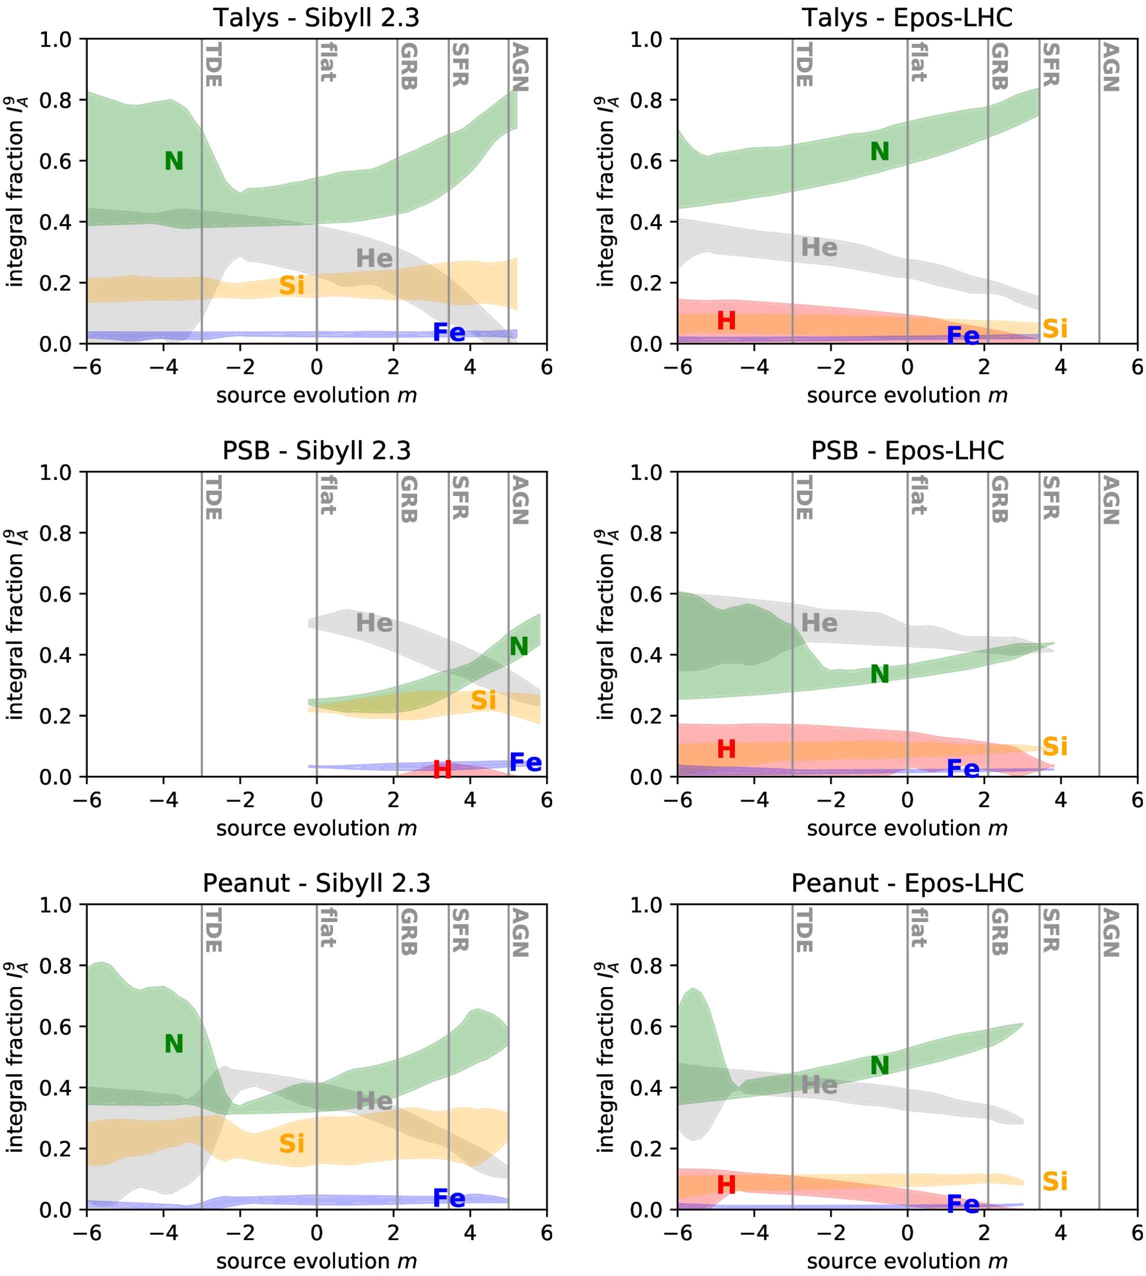

Figure 5.7: Ranges in the fraction allowed within 3σ (for 2 d.o.f) as a function of the source evolution parameter. The fractions are defined by integrating the injection spectrum from $E_{\text{min}} = 10^9$ GeV, see Eq. (5.3). The fractions are given for all combinations of Talys, PSB and Peanut with Sibyll 2.3 and Epos-LHC. QGSjetII-04 is omitted as it does not provide a good fit to the observables. Figure extended from Heinze et al. [3]

Within the limitations of our generic model, we illustrate the ranges of the injected fractions I_A^9 within the 3σ contours of our fit in Fig. 5.7 as a function of the source evolution. The figure shows the different combinations of Talys, PSB and Peanut with Sibyll 2.3 and Epos-LHC. This corresponds to the left and middle column in Fig. 5.4. As QGSjetII-04 does not give a consistent interpretation of the composition data, we leave it out of the discussion here.

Comparing the fraction ranges for Sibyll 2.3 (Fig. 5.7, left column) with respect to Epos-LHC (Fig. 5.7, right column), the most striking difference is in the silicon fraction, which is significantly higher for Sibyll 2.3, while in turn the nitrogen fraction is higher for Epos-LHC. This is mainly due to the heavier $\langle \ln A \rangle$ predicted by Sibyll 2.3. The Epos-LHC model allows for a significant fraction of proton injection with $\lesssim 15$ %. The main reason for this is the lighter $\langle \ln A \rangle$. However, this is also affected by the slightly lower rigidity found for this model, which makes the protons pile up mostly below the fit range. This is illustrated by the best fit for the combination of Epos-LHC and Peanut, shown in the appendix in Fig. B.5 and Fig. B.6. For both hadronic models, the nitrogen fraction increases at the cost of the helium fraction with higher source evolution. The higher disintegration for distant sources produces more helium during propagation, therefore requiring less helium injected at the source.

The combination of $\langle X_{\max} \rangle$ and $\sigma(X_{\max})$ predicted by Sibyll 2.3 allows for a stronger superposition of different mass spectra. Therefore, Sibyll 2.3 leaves the mass fractions less constrained compared to Epos-LHC for the same source evolution. In both cases, the allowed mass fractions widen when going to negative source evolution. This is an effect of the correlation in $\gamma - m$. Strong source evolution is found together with hard spectral indices, i.e. rather peaked spectra. The exact shape of the spectrum then has little influence on the integral fraction I_A^9. At negative source evolution, the spectral indices are softer, such that lower energies contribute more to I_A^9. For the same source evolution, the exact value of γ then has a stronger influence on I_A^9. This effect is illustrated by the two examples of injection spectra for the baseline case in Fig. 5.2 and Fig. 5.3.

The choice of the disintegration model affects mainly the fraction of injected He and Ni. As mentioned in Section 5.3.2, negative source evolution is not contained in the 3σ contours for the combination of PSB and Sibyll 2.3. This constrains the fraction ranges in Fig. 5.7 (middle left panel) to positive source evolution. The most relevant features of the disintegration model are the level of α emission and the number of open reaction channels that control how efficiently the nuclear cascade develops, as discussed in detail in Section 3.2.5. For instance, the absence of α emission in PSB, is compensated by higher He fractions at the source, as noticed in Aab et al. [148], Alves Batista et al. [210]. Due to the less efficient photo-disintegration in PSB, the nuclear cascade can only develop sufficiently for more distant sources. In turn local sources are disfavoured when using this model. A similar effect can be seen by comparing Peanut and

Talys, though the difference is less pronounced and negative source evolution is not excluded for either model. As in the Talys model α-emission is most efficient, it requires the least He injection at the source.

The fractions shown in Fig. 5.7, describe the the integral fraction ejected from the source. They have important implications on the physics inside the source. Disintegration inside the source produces more protons and helium. Therefore higher-luminosity sources are more likely to eject a large fraction of these lighter elements. For strong evolutions, such as for AGNs, the heavier isotopes must escape mostly intact. Weaker source evolutions seems to allow for higher helium and maybe even proton fractions, which allows for partial disintegration inside the source. This however gives only a rough estimate of the required composition. The details need to be studied in the context of a specific source model.

Note that we find a small, but non-zero iron fraction throughout all model combinations. This is a result of the increased statistics at the cutoff in the updated Auger 2017 data set in combination with the rigidity-dependent maximal energy, as discussed in Appendix B.1.

5.4 Cosmogenic neutrino fluxes

With the source parameters fixed by the fit to the UHECR data, a prediction of the cosmogenic neutrino flux is also implied. However the UHECR fit is only sensitive up to a redshift of $z_{\max} = 1$. Contrariwise, cosmogenic neutrinos can come from a redshifts above one, as they are not absorbed by interactions. Their flux is only reduced by the adiabatic expansion of the universe. Therefore, it is impossible to estimate a robust confidence interval for cosmogenic neutrinos using a method solely based on UHECR data.

To get a conservative estimate of the cosmogenic neutrino flux, we consider only the redshift range to which the UHECR fit is sensitive, i.e. $z_{\max} = 1$. In Fig. 5.8, we show the neutrino ranges corresponding to the 1σ, 2σ and 3σ contours of the fit with the baseline model combination (Talys-Sibyll 2.3). The parameter space in $\gamma - m$ is shown again for comparison. The dominant impact on the cosmogenic neutrino flux is due to source evolution, which can be seen by comparing the ranges to the contours in parameter space. In contrast to the 1σ region, which is limited to positive source evolutions, the 3σ region is unconstrained towards negative evolution. Hence, if the sources are local, the expected cosmogenic fluxes are very low. The peak energy of the flux at $\sim 10^8$ GeV implies that these neutrinos are mainly produced on the EBL. This is because the maximal rigidity found in the fit is not high enough to reach the threshold for pion-production on the peak of CMB. For higher maximal rigidity the flux from interactions with the CMB would peak at $\sim 5 \cdot 10^9$ GeV and be much higher due to the higher number density of the CMB [5].

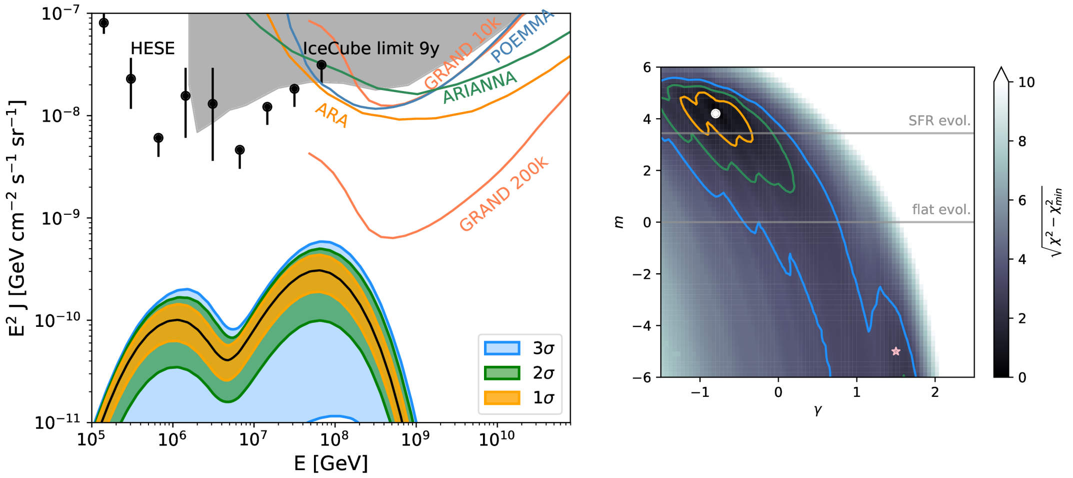

Figure 5.8: Allowed range of the neutrino flux (all flavors) from the 3D fit in Fig. 5.4 within the $1\sigma, 2\sigma, 3\sigma$ contours (for 2 d.o.f.). The source evolution is in this case defined as $(1+z)^m$ for a maximum redshift $z_{\max} = 1$, where m changes within the allowed fit regions. Estimated sensitivities for future radio neutrino detectors are shown for comparison: ARA [235], ARIANNA [236], GRAND [103] and POEMMA [237]. The small panel shows the corresponding contour in $\gamma - m$ for comparison(this is the same as the lower left panel in Fig. 5.4). Figures taken from Heinze et al. [3].

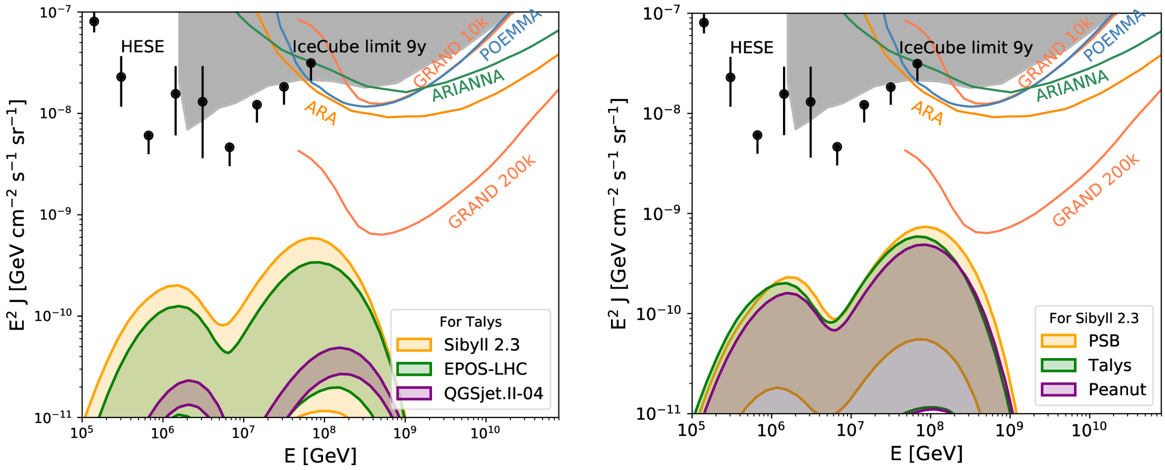

Figure 5.9: Allowed range for the neutrino flux (all flavors) in the 3σ region for different model assumptions. Left: The disintegration model is fixed to Talys and the ranges for different air-shower models are shown. Right: The shower model is fixed to Sibyll 2.3 and the ranges for different disintegration models are shown. Figure taken from Heinze et al. [3].

As the maximum rigidity is strongly constrained by the UHECR fit, the high-energy peak of the neutrino flux stays relatively robust and located at $\sim 10^8$ GeV. This is in agreement with Alves Batista et al. [27], where equally low fluxes were predicted. A small but relevant difference resides in the propagation code, since Alves Batista et al. [27] assume a simplified redshift scaling of the EBL. The effects of this scaling on the neutrino fluxes are explained in Alves Batista et al. [153]. If we apply the same simplified scaling, the cosmogenic neutrino flux in our calculations increases by 50%. Other minor differences come from other details of the propagation code and the fitting procedure. Differences to other works come from their limiting assumptions about the source evolution, injected composition or the cutoff energy. Romero-Wolf and Ave [151] fit UHECRs with a bayesian method, but restrict their fit to positive source evolution ($m \in [3.4, 5]$), leading to high neutrino predictions. Møller et al. [154] consider only two injected elements (proton and iron) with the maximal energy fixed to $2.5 \cdot 10^{11}$ GeV. For iron this corresponds to a rigidity $R_{\text{max}} = 10^{10}$ GV, which is almost an order of magnitude higher than in our best fit. This choice again leads to higher neutrino predictions. Das et al. [152] discuss a fit with only proton and helium injection as their main result, again leading to high neutrino fluxes due to the higher required rigidity. For a mixed composition, they find similar a neutrino flux level as in our results.

We now focus on the 3σ contours to study the robustness of our results against changes of the disintegration and the air-shower model. The ranges for different model combinations are shown in Fig. 5.9. In the left and right panel of Fig. 5.9 the cosmogenic neutrino flux is shown corresponding to the blue UHECR contours for the models in the top row and left column of Fig. 5.5, respectively. As the neutrino flux is mainly impacted by the source evolution, the main difference between the models therefore comes from the allowed range in m. For the Sophia-superposition model described in Section 3.2.5, the neutrino production rate depends on the energy per nucleon which is similar for all species due to the fixed rigidity. Hence the composition of the neutrino flux dependence is weak for our generic source model. The variations between the disintegration models are small, resulting in a relatively robust upper bound. For QGSjetII-04 the flux is small since positive evolutions are disfavoured. For PSB, a sizable lower limit to the neutrino flux exists, since negative source evolution (local sources) is not contained in the fit ranges.

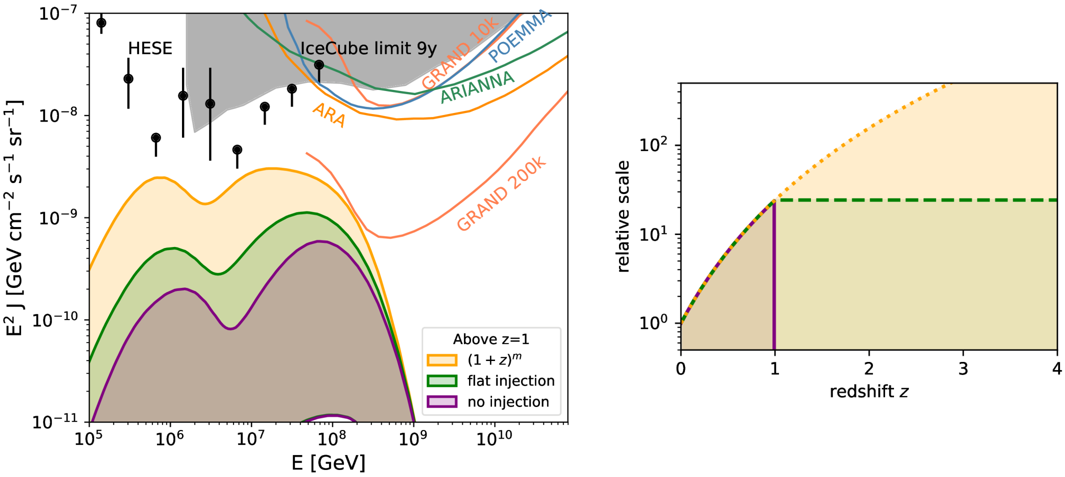

Figure 5.10: Allowed range for the neutrino flux (all flavors) in the 3σ region for different source evolution. The purple range corresponds to $z_{\max} = 1$ (same as Fig. 5.8). For the other curves the source evolution is continued to $z_{\max} = 3$ either by continuing as $(1+z)^m$ (yellow) or with a break to flat evolution at $z = 1$ (green). The smaller panel illustrates the different choices for the source evolution. Left panel taken from Heinze et al. [3].

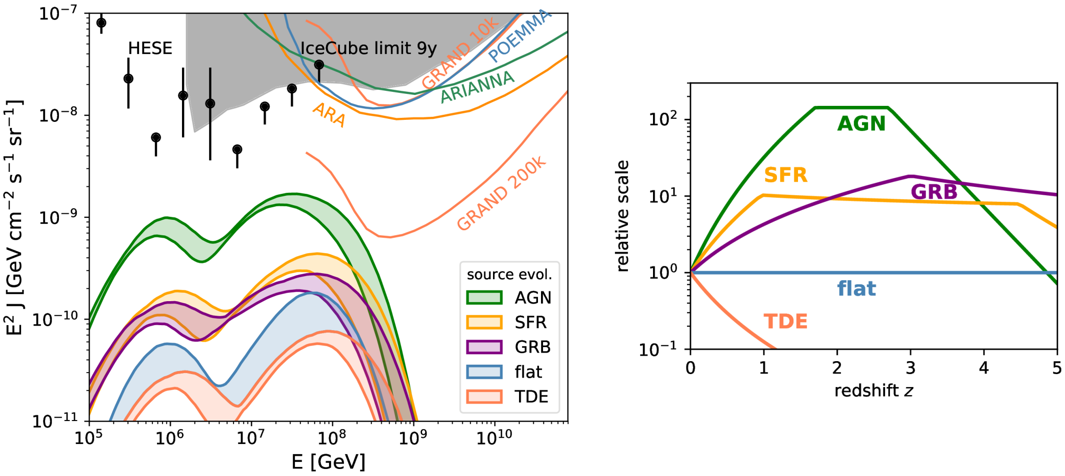

Figure 5.11: Allowed range for the neutrino flux (all flavors) in the 3σ region for the source evolution fixed to different source classes and for flat evolution. The smaller panel shows the different source over redshift. Left panel taken from Heinze et al. [3].

The most significant impact on the neutrino fluxes comes from the extrapolation of the source evolution to redshifts $z > 1$, which is unconstrained by UHECR data. We adapt two approaches for the extrapolation:

An empirical method using a simple continuation of the $(1 + z)^m$ parameterization beyond $z = 1$ up to $z_{\max} = 3$. We also test a distribution with a break at $z = 1$ and a flat $(m = 0)$ behavior beyond that. The source evolutions and resulting fluxes are shown in Fig. 5.10

Discrete evolution functions of candidate source classes, where the parameter m is not free; AGN [238, 239], GRB [240], SFR [241] (this evolution includes starburst galaxies), TDE [242] and a flat evolution. In this case $z_{\max} = 5$ is used, which is above the cutoff for all source evolutions used. The source evolutions and resulting fluxes are shown in Fig. 5.11

The most optimistic $(1 + z)^m$ extrapolation results in fluxes that are one order of magnitude below the diffuse neutrino flux observed by IceCube. Since this choice is rather extreme, it can be considered as the upper limit on the flux of cosmogenic neutrinos for a single dominant UHECR source population with a rigidity-dependent energy cutoff. A flux at a similar level is found for AGN evolution. In either scenario, the planned radio-based neutrino detectors will neither be able to distinguish between source types (Fig. 5.11) nor detect any significant cosmogenic neutrino signal. It is important to note that a lower-limit on the cosmogenic neutrino flux can only be defined in a meaningful way, if the source evolution is constrained, either by the choice of a specific source model or by more information on the UHECR sources from future observations.

5.5 Summary and discussion

In this chapter, we presented a fit of the 2017 Auger energy spectrum and composition measurements. For this, we assumed a single dominant source of UHECRs that accelerates nuclei up to iron with a rigidity dependent cutoff. Compared to the Auger Combined Fit (CF) [148], we extended this framework to include the source evolution as an additional free parameter. Within this model, we studied the effect of the main model uncertainties affecting the UHECR composition at the source: Nuclear disintegration during propagation and hadronic interactions in the air-shower development.

In the three-dimensional scan over the main source parameters, $(\gamma, R_{\max}, m)$, we find a low maximal rigidity for all parameter combinations, implying that the observed UHECR cutoff is not due to the GZK-effect, but instead due to the sources reaching their maximal acceleration. In $\gamma - m$ we see a clear correlation: Sources with softer spectral indices $(\gamma \gtrsim 1)$, that are compatible with diffusive shock acceleration, must be local $(m < 0)$. Source classes satisfying this would be

Tidal Disruption Events [129–131] and low luminosity GRBs [133, 134]. Instead, sources with stronger source evolution (like star forming rate) require very hard, almost monochromatic, spectral indices, which are more difficult to explain with a standard acceleration mechanism. Possible examples are re-acceleration scenarios as those proposed for termination shocks in starburst and nearby radio galaxies [127, 243, 244]. Both the low maximal rigidity and the $\gamma - m$-correlation are driven by the X_{max}-distribution, which allows only for little overlap of different masses. A larger maximal rigidity leads to stronger disintegration of primaries, with the secondary nuclei increasing the overlap of different masses. The same is true for a large population of distant sources, corresponding to large values of m in our fit. This is partly compensated by the hard spectral indices, which decreases the overlap of different mass spectra.

We discussed the impact of model variations for or all combinations of the disintegration models PSB, PEANUT, and TALYS and the air-shower models EPOS-LHC, SIBYLL 2.3, and QGSJETII-04. The hadronic interaction model has the largest effect on the UHECR fit. Depending on the hadronic model the fit either favors local sources (for EPOS-LHC) or those with stronger evolution (for SIBYLL 2.3). Still, the 3σ-region in both cases contains most of the range in m, such that this result alone does not give a strong conclusion about a preferred source type. While the source evolution is not constrained from this generic fit, the ejected mass composition is limited, preferring a mix of mostly nitrogen and helium. The fraction of silicon and hydrogen depends mostly on the choice of hadronic interaction model, while the disintegration models affect the fraction of helium injection. Almost independently of the model combinations, we find a small but non-zero iron fraction in the injection. This result is dependent on our choice of rigidity-dependent maximal energies: In combination with the low maximal rigidity, that is fixed by the fit to the *ankle* region, iron injection is needed to reach the UHECR cutoff. However, we also showed by a more general injection function that a rigidity-dependent injection is strongly preferred by the observed X_{max}.

More generally, these results show the limitations of what can be inferred with the current statistics of UHECR data. The data is well described by a generic, rigidity-dependent source model, which qualitatively represents a wide range of possible source classes. Still, a strong degeneracy in the parameter space remains, such that no single source class is strongly preferred. Including data at lower energies below the *ankle* might break this degeneracy. However, this requires additional assumptions about extragalactic magnetic field and the transition to a (possibly) Galactic component below the ankle. These assumptions would again add a high degree of uncertainty to the model. We did not include the effects of magnetic diffusion in extragalactic magnetic fields in the propagation study. This assumption is only strictly valid for a continuous source distribution [201]. The impact of magnetic diffusion depends on the separation between individual sources and the strength of intergalactic magnetic fields. It therefore introduces several additional parameters with high uncertainties. It was found in Wittkowski [149] that the inclusion of magnetic fields leads to somewhat softer γ, as the propagated spectra at lower energies are suppressed. However, apart from shifting the range in γ, we do not expect magnetic diffusion to affect our results qualitatively.

Additionally, there is an ambiguity in the exact composition due to the large range of mass fractions at the source allowed by the different air-shower models. With new data from future experiments the parameter space will likely be more constrained. For example, with more accurate data on the UHECR composition from the AugerPrime upgrade [85], the injected mass fraction will likely be more constrained. This improved sensitivity might however be sufficient to break the ambiguity between the different photo-disintegration and air-shower models. Significant progress on these models is therefore needed as well to make robust conclusion about the sources of UHECRs.

The allowed parameter space for the sources of UHECRs could in principle be constrained by secondary messengers such as cosmogenic neutrinos, if future experiments were sensitive enough to observe them. We therefore computed the cosmogenic neutrino flux resulting from our fit to the UHECR data. A lower bound on the cosmogenic neutrino flux is not meaningful, as local source evolution is not constrained by the fit. However, we find that the upper bound on the cosmogenic neutrino flux is relatively robust even under model variations. The expected low flux level makes the detection in future radio based detectors very unlikely. This result is valid for the assumption of a generic and dominant rigidity dependent source class, which is preferred by the X_{max} data that allows only for a small overlap of masses. However, if a subdominant population of sources would accelerate protons to higher energies, they would produce neutrinos more efficiently off the CMB and could enhance the cosmogenic neutrino flux significantly [154, 155, 159]. On the positive side, the low expected flux of cosmogenic neutrinos leaves room for the detection of source neutrinos in the future, if there are sources able to produce neutrinos at energies beyond the flux observed by IceCube.

Chapter 6

Neutrinos and UHECRs in multi-collision models of GRBs

In this chapter, we discuss gamma-ray bursts (GRBs) in a multi collision internal shock model. Within this model, the relativistic plasma outflow from a GRB is approximated by a sequence of discrete shells that collide at different radii producing shocks that accelerate particles. We study the impact of two model assumptions on the multi-messenger output of the GRB: The initial distribution of shells and the hydrodynamic model of the collision between a pair of shells. For these assumptions, we predict the expected flux of source neutrinos and discuss its detectability. We also discuss the UHECR output spectra qualitatively in relation to the parameters required by the UHECR fit discussed in the previous Chapter 5.

This chapter is based on the previous publications in Bustamante et al. [1] and Rudolph et al. [4]. I reimplemented the underlying multi-collision in PYTHON from existing MATHEMATICA scripts to allow for easier model variations. I also contributed to updating the radiation code based on NEUCOSMA to allow for the treatment of UHECR nuclei. For the first study, my main contribution was analysing and adjusting the example cases after qualitative changes in the updated model. Since there were several updates to the model since the initial publication, the results from this study are updated for this chapter using the most recent version of the model. For the second study, I contributed to cross-checking the implementation of the alternative collision models and evaluating the results.

6.1 Gamma-ray bursts

Gamma-ray bursts (GRBs) are among the most luminous astrophysical objects observed [245, 246]. They are short outbursts of gamma rays with energies up to 100 GeV that last only a few tens of seconds. With luminosities of $10^{47} - 10^{53}$ ergs s^{-1} and a local rate of 1.3 Gpc^{-3} yr^{-1} they are in principle abundant and luminous enough to power UHERCs. They were therefore long believed to be prime candidates for the sources of UHECRs. However, this theory has recently been challenged due to the lack of neutrino observations in coincidence with GRBs [113].

The *Burst and Transient Source Explorer* (BATSE), operating from 1991, detected more than 2700 GRBs over its operation time [247], with different luminosities, durations and light curves. The distribution of the observed GRBs is highly isotropic [248]. Some examples of light curves detected by BATSE are shown in Fig. 6.1. Another detector is the *Swift Gamma-Ray Burst Mission* (SWIFT), that was launched in 2004 [30] and is still operating today. *The Fermi Gamma-Ray Space Telescope* (FERMI) carries two other important gamma-ray-burst detectors, the *Gamma-Ray Burst Monitor* (GBM) [32] and the *Large Area Telescope* (LAT) [31].

The observed gamma-ray bursts can be divided into two classes by their duration. Long GRBs (LGRB) last between 10 and 100 seconds, while shorts GRBs have durations of less than 2 seconds [249]. LGRBs have been observed in coincidence with core-collapse supernovae [250], whereas SGRBs are likely to originate from mergers of compact objects (neutron stars or black holes) [251]. The latter hypothesis was recently confirmed by the coincident observation of GRB170817A with the binary neutron star merger GW170817 [42]. About 84 % of the GRBs observed with FERMI-GBM are long GRBs [252]. For both classes of GRBs, the observed electromagnetic radiation is believed to originate from relativistic plasma jets that are accelerated by the central engine.

Their high luminosity and inferred high particle densities make GRBs prime candidates for the sources of UHECRs. However, a significant abundance of baryons at the acceleration zone in combination with high luminosities, should lead to an efficient production of astrophysical neutrinos. Additionally, their short duration makes them ideal for neutrino detection, as the short time window significantly reduces backgrounds. Still, the IceCube collaboration has found no significant coincidence of astrophysical neutrinos with high-luminosity GRBs, which puts stringent stacking limit on their contribution to the astrophysical neutrino flux [25]. The simpler models of the UHECR-GRB connection are therefore mostly excluded, in which cosmic ray acceleration and neutrino production happens in a single zone in the GRB jet [113, 124, 224, 253]. However, if the production volume is optically thin to photo-hadronic interactions, the neutrino production is reduced. This is for example the case for multi-collision models, which consider several regions in the GRB jet. Neutrino and UHECR production is then dominated by different radii, which relaxes the stacking bounds [254]. Another way to avoid the stacking limits is to consider only low-luminosity GRBs (LL-GRBs) with luminosities below 10^{49} erg s^{-1}. They are more abundant and have longer durations than conventional GRBs, and are partly below the sensitivity of current experiments due to the lower luminosity. This makes neutrino limits less constraining. Therefore, LL-GRBs have been considered as the sources of both UHECRs and astrophysical neutrinos in the more recent literature. [132–134, 255, 256].

The electromagnetic radiation of GRBs shows a strong time variability, which can be as low as 10 ms. Several light curves observed by BATSE are shown in Fig. 6.1. The observed energy

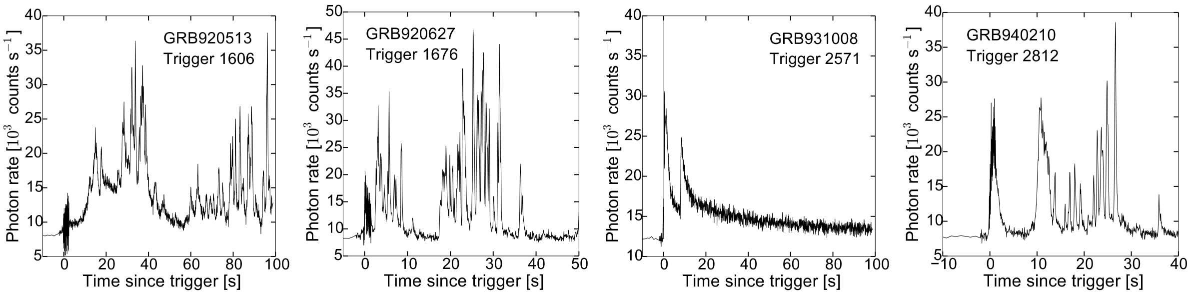

Figure 6.1: Four selected light curves measured by BATSE [257] in the $>$ 20keV range. Figures taken from Bustamante et al. [1].

spectra of GRBs are non-thermal and different models for the exact spectral shape exist. In their standard analysis the Fermi-GBM collaborations fits four different models: a simple power law, Band's GRB function, an exponential cut-off power-law (comptonized model) and a smoothly broken power law [252]. The mechanism to generate the spectrum depends strongly on the assumed model, such as whether a purely leptonic mechanism or an abundance of hadrons is assumed. An additional observable is the GRB afterglow, radiation in the x-ray and/or optical range that can be observed after the initial prompt gamma-ray emission. The standard interpretation is that the afterglow is produced by the interaction of the relativistic plasma with the surrounding medium [258].

Different models for the dissipation mechanism of GRBs exist, see Pe'er [259] for a detailed review. In the internal shock model, the radiation is created by internal collisions of the plasma outflow, which generate relativistic shocks that dissipate a fraction of the internal kinetic energy as radiation [260–262]. Non-thermal electrons that are accelerated in the shock then produce the prompt gamma-ray emission through synchrotron radiation. The afterglow is generated by the interaction of the bulk outflow with the surrounding medium. This requires a large efficiency of the internal shocks, as otherwise the observed afterglow would be too luminous compared to the prompt emission, which is inconsistent with observations [263, 264].

In photospheric models [265–268], most of the energy is dissipated while the jet is optically thick to photons (i.e. below the photosphere) resulting in a quasi-thermal spectrum, which is then upscattered by e.g. inverse Compton-scattering to produce non-thermal radiation. Magnetic reconnection models are another alternative [269–274]. In these models, a large fraction of the dissipated energy is initially transferred to magnetic fields which then dissipates energy at larger radii through magnetic reconnection. Even in these two alternative models, internal shocks can play an important role in transferring the initial kinetic energy to radiation. For example the sub-photospheric emission in photospheric models might be driven by internal shocks [265]. The model presented in this chapter is a version of the classical internal shock model.

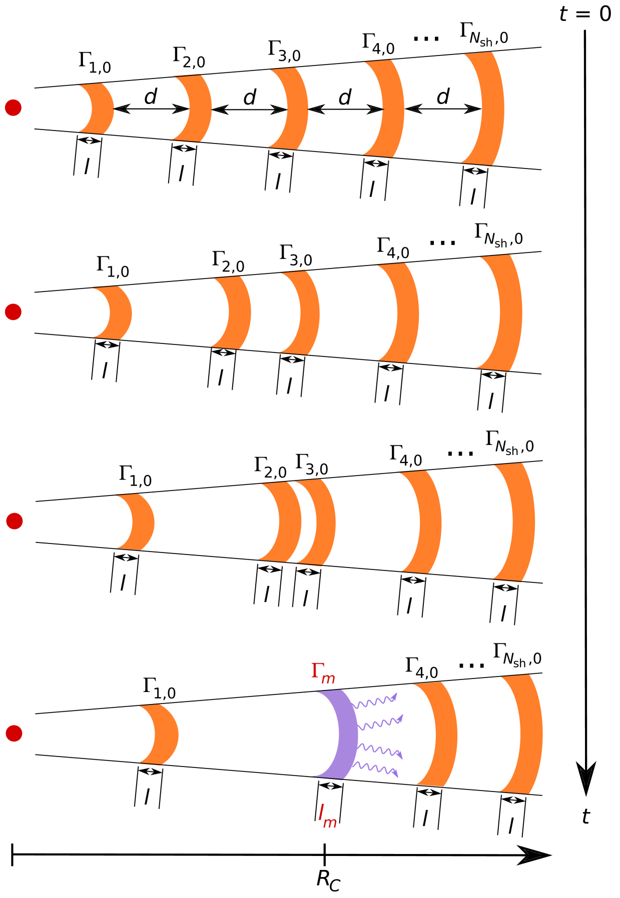

Figure 6.2: Schematic illustration of the fireball evolution from the initialization with equal distances, evolution and collision of a pair of plasma shells. See Section 6.2.1 for details. Figure taken from Bustamante et al. [1].

6.2 Internal shock model with multiple collisions

In the internal shock model, the observed GRB emission is generated by an outflow of plasma shells. When they collide with each other, (mildly) relativistic shocks form that accelerate particles to relativistic speeds, converting the kinetic energy of the shells to radiation. A schematic illustration is shown in Fig. 6.2.

The properties of the light curve like its duration T, luminosity L_γ and variability timescale t_ν are then attributed to distribution of the parameters of the initial shells, such as the separation, Lorentz factor and kinetic energy. Depending on the number and distribution of initial shells, each shell can collide and merge multiple times during the fireball evolution. This results in a distribution of collisions with different parameters, which depend non-linearly on the initial engine setup. For a stochastic distribution of the initial shell parameters, the collision parameters will be distributed around an average collision radius, where most of the energy is dissipated. One-zone collision models therefore assume only a single collision, for which representative parameters are derived from observations, such as the bulk Lorentz factor $\langle\Gamma\rangle$ and average luminosity [121, 124, 224, 275–278]. Assuming that the variability in the observed light curve t_ν is connected to the initial separation between shells leads to relations for the collision radius R_{coll} and width of the emitting region l_{coll}:

$$R_{\mathrm{coll}} = 2\langle\Gamma\rangle^2 \frac{c \cdot t_\nu}{1+z} \quad \text{and} \quad l_{\mathrm{coll}} = \langle\Gamma\rangle \frac{c \cdot t_\nu}{1+z} \,. \tag{6.1}$$

The relation for the radius R_{coll} is derived under the assumption that the average relative velocity between shells is $\Delta\beta \approx 1/2\Gamma^2$ [275, 279]. The same relations can also be derived independently from the curvature of the emitting shell [280]. The results obtained from this single collision are then scaled up to the whole burst assuming a total number of collision $N_{\mathrm{coll}} = T_{90}/t_\nu$, where T_{90} is the time in which 90% of the energy is dissipated.

This simplified one-zone model is justified given that observations only provide averaged estimates such as a bulk Lorentz factor $\langle\Gamma\rangle$ and total luminosity L_γ. Still, in more realistic models, the collisions are distributed over a wide range of radii, reaching from below the photosphere to the circumburst medium. This affects the particle output, as the acceleration and production efficiencies for different messengers scale differently with the radius. Multi-Collision models therefore simulate the full evolution of the GRB fireball from its initial setup to the circumburst medium. These models then also predict the time profile of the light curve by assuming that each collision emits a gamma-ray pulse.

In this chapter, we discuss such a multi-collision model for long GRBs. We assume a full distribution of $N_{\mathrm{sh}} = 1000$ initial shells, which evolve dynamically in the fireball by colliding and merging. The radiated spectra are then calculated as the superposition of all individual collisions. This approach introduces many new free model parameters such as the initial shell distribution and the hydrodynamic model used for individual collisions. The discussion of these models parameters is complicated by the fact that they can only be indirectly related to the observations such as light curves and gamma-ray spectra. We therefore focus mostly on the qualitative features and their impact on the prompt neutrino fluxes.

Multi-collision models have been developed to explain fast time variability of GRBs and the variety in their observed light curves. Instead of calculating the outflow from GRBs in a full magneto-hydrodynamic simulation, these models approximate it with a discrete number of shells and make use of simple analytic approximations for the collision of a pair of shells [281–285]. While some of these models make slightly different assumptions in the analytic descriptions, a major difference comes from the initial setup of the shells. The initial shells are emitted with a distribution of Lorentz factors Γ, which is in principle arbitrary. The distribution should reproduce observed GRB light curves qualitatively, however it is not straightforward to predict from the initial Lorentz factor distribution without a full simulation. Similarly, the energy and mass distribution (related by $E_{\text{kin}} = \Gamma \cdot m$) is a free parameter.

From the initial setup, the fireball is evolved over all collisions between two shell, which in most models merges them into a single shell. Each collision transfers a fraction of the initial kinetic energy into internal energy that is radiated away. The evolution then proceeds until the shells reach the circumburst medium. An important quantity is the efficiency of conversion of kinetic to internal energy $\epsilon = E_{\text{diss}}/E_{\text{kin}}$. It depends both on the initial distribution of shells as well as the details of the collision mechanics. The remaining kinetic energy will be visible in the afterglow, for which the energy budget is constrained by observations.

We use the model from Kobayashi et al. [281] as our collision model and explore the impact of different initial shell distributions and collision assumptions. The collision model is introduced in more detail in Section 6.2.1. To calculate the radiation from the individual collisions, we use the NEUCOSMA code. The details of the radiation model are given in Section 6.2.2. In practice, the fireball evolution is computed in an object-oriented PYTHON code, which were fully reimplemented from existing MATHEMATICA scripts originally written for Bustamante et al. [254]. This code generates the relevant collision parameters for the radiation calculations. These are then handed over to a C-code based on NEUCOSMA to perform the performance intensive radiation calculations. The version used in Bustamante et al. [1] was only able to treat protons. We since updated the C-code to the most recent NEUCOSMA version to treat the full nuclear cascade, while also significantly adjusting the code and interface to PYTHON to allow for more flexibility in changing the model assumptions.

6.2.1 Hydrodynamic collision model

A GRB jet is a continuous outflow with fluctuating energy and particle densities. We refer to the regions with high densities as shells. A full magneto-hydrodynamic simulation even in one dimension is computationally expensive. We therefore use a simplified analytic approximation following Kobayashi et al. [281], in which each shell is box-shaped with constant width and density.

In this simplified model, the k-th shell is fully defined by its mass m_k (or kinetic energy $E_{\text{kin},k}$), width l_k, Lorentz factor Γ_k and (current) radius R_k. Other dependent properties can be derived, like the speed β_k, the volume $V_k = 4\pi R_k^2 l_k$ (assuming spherical geometry) and density ρ_k. In principle each property of each individual shell is a free parameter of the model. However, the shell parameters are usually defined from a stochastic distribution with a few macroscopic burst parameters. We will discuss both stochastic and structured initial distribution in Section 6.4. The choice for the initial kinetic energy $E_{\mathrm{kin},k}$ of each shell is another free parameter, typical choices are equal-mass or equal-energy. We will discuss the impact of this choice later. From this setup, the system is evolved until all shells reach the circumburst medium at R_{max}, giving a distribution of collisions with different parameters. Note that the one-zone estimate in Eq. (6.1) is derived using a stochastic initial distribution in Γ_k, for more structured distributions in Γ_k it does not necessarily hold.

A collision happens when a *rapid* (index r) catches up to a *slow* (index s) shell. The *merged* (index m) shell's Lorentz factor Γ_m is derived from momentum conservation and assuming relativistic shells ($\beta_k \simeq 1 - 1/(2\Gamma_k^2)$), resulting in

$$\Gamma_m = \sqrt{\frac{m_r\Gamma_r + m_s\Gamma_s}{m_r/\Gamma_r + m_s/\Gamma_s}}\,. \tag{6.2}$$

The internal energy available for radiation is then the difference of kinetic energies before and after the (inelastic) collision:

$$E_{\mathrm{int},m} = m_r\Gamma_r + m_s\Gamma_s - (m_r + m_s)\Gamma_m\,. \tag{6.3}$$

The simplest assumption is that all of the internal energy is radiated in relativistic particles ($E_{\mathrm{diss},m} = \eta\, E_{\mathrm{int},m}$ with $\eta = 1$). Modifications to this assumption will be discussed in Section 6.5.

The Lorentz factors of the forward (fs) and the reverse shock (rs) are determined by

$$\Gamma_{\mathrm{fs\,(rs)}} = \Gamma_m \cdot \sqrt{\frac{1 + 2\Gamma_m/\Gamma_{s\,(r)}}{2 + \Gamma_m/\Gamma_{s\,(r)}}}\,. \tag{6.4}$$

The timescale of emission t_{em} is estimated from the time that the reverse shock takes to cross the rapid shell

$$\delta t_{\mathrm{em}} = \frac{l_r}{c\,(\beta_r - \beta_{\mathrm{rs}})}\,. \tag{6.5}$$

Note that this timescale is observed Doppler-boosted. The merged shell is compressed by the shocks, its width is computed as

$$l_m = l_s\frac{\beta_{\mathrm{fs}} - \beta_m}{\beta_{\mathrm{fs}} - \beta_s} + l_r\frac{\beta_m - \beta_{\mathrm{rs}}}{\beta_r - \beta_{\mathrm{rs}}}\,. \tag{6.6}$$

The parameters of the collision (E_{diss}, $\Gamma_{\mathrm{coll}} = \Gamma_m$, δt_{em}, $l_{\mathrm{coll}} = l_m$, $R_{\mathrm{coll}} = R_m$) are then used for the radiation modelling. The merged shell then continues to propagate in the fireball.

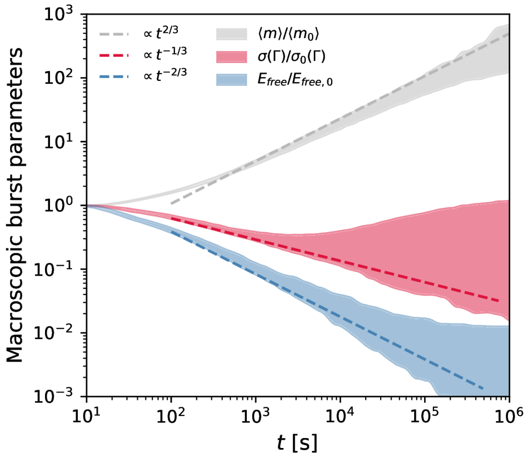

Figure 6.3: Time evolution in the source frame of average shell mass $\langle m\rangle/\langle m_0\rangle$, standard deviation of the Lorentz factor $\sigma(\Gamma)/\sigma(\Gamma_0)$, and total free energy of a burst, $E^{\rm iso}_{\rm int,tot}/E^{\rm iso}_{\rm int,tot,0}$. The ranges are from numerical results for 1000 simulations run with random initial setup for the parameter values $N_{\rm sh}=3000$, $\Gamma_0=100$, $A_\Gamma=0.2$, $\delta t_{\rm eng}=10^{-3}$ s, $d=l$, $z=2$, and $E^{\rm iso}_{\rm kin,0}=10^{52}$ erg. The dashed lines show analytical estimates for the time evolution from from Beloborodov [283]. Figure updated from Bustamante et al. [1].

Due to momentum conservation there is a limit on the total energy that the fireball can dissipate as radiation. It is reached when all shells have merged into a single shell with the total mass $M_{\rm tot}=\sum_k m_k$. The available free energy of a burst can therefore be calculated by the difference in kinetic energy of the individual shells to the fully merged fireball:

$$E_{\rm free}=\sum_k(\Gamma_k\cdot m_k-\Gamma_{\rm cm}\cdot m_k)=\Gamma_{\rm cm}(M_{\rm inv}-M_{\rm tot})\,, \tag{6.7}$$

where $M_{\rm inv}=\sqrt{E^2_{\rm tot}-p^2_{\rm tot}}$ is the invariant mass of the system and the center of mass Lorentz factor is calculated as $\Gamma_{\rm CM}=E_{\rm kin,tot}/M_{\rm inv}$. The available free energy is therefore larger for higher spread in the initial Γ factors. Since in each collision a fraction of the available kinetic energy is radiated away, the spread in Γ decreases with the fireball evolution. As a result of mass conservation, the average mass per shell increases with each collision.

This evolution is illustrated in Fig. 6.3, which shows the time evolution (in the source frame) of macroscopic burst parameters for a full fireball simulation: Average shell mass $\langle m\rangle/\langle m_0\rangle$ (subscripts of zero indicates values at simulation start), standard deviation of the Lorentz factor $\sigma(\Gamma)/\sigma(\Gamma_0)$, and total available internal energy of the burst $E^{\rm iso}_{\rm int,tot}/E^{\rm iso}_{\rm int,tot,0}$. The numerical results of our simulation match the analytical power-law estimates from Beloborodov [283], which assume small fluctuations in the initial Lorentz factors, i.e. $A_\Gamma \ll 1$. Deviations are visible at late times, when the number of remaining shells is low and the analytical predictions are no longer applicable. This late deviation depends strongly on the random initial setup, so we show ranges obtained after running 1000 different simulations.

For the calculations of the shell and the collision properties, we follow exactly the model presented in Kobayashi et al. [281]. Some of the analytic formulas are slightly different in Daigne and Mochkovitch [282] (also used e.g. in Globus et al. [122], Bosnjak et al. [286]). For example the merged Lorentz factor Γ_m in Eq. (6.2) is the value after the two shells have fully merged. If the shells have significantly different masses, most of the energy is dissipated in the initial phase of the collision, where the lighter shells sweeps up a mass equal to its own in the more massive shell. Daigne and Mochkovitch [282] therefore use Eq. (6.2) for equal mass as the more representative formula to calculate the Lorentz factor during the energy dissipation $\Gamma_{\text{coll}} = \sqrt{\Gamma_r \Gamma_s}$. For the merged shell and dissipated energy one still has to use Eq. (6.2) and Eq. (6.3) to comply with energy and momentum conservation. Also, Daigne and Mochkovitch [282] use a different estimate for the energy density in the radiation zone, without explicitly relying on the merged shell width l_m. Note that the degeneracy between these model choices cannot be broken, as neither is fully accurate. They all approximate a continuous outflow, where neither of the parameters are actually discrete. Important for the results is the qualitative scaling of the collision parameters, such as its volume, which scales with R^2_{coll} in both cases. While some quantitative differences therefore may come from the model choices, the qualitative features are the same in either model.

6.2.2 Radiation model

The collision parameters derived in the previous section are used for the computation of the energy density and the emission spectra in the merged shell frame (primed). The computations are done in the isotropic-equivalent picture assuming an observation angle of $\theta_{\text{obs}} = 1/\Gamma$. The total luminosity injected into the radiation zone is $L'_{\text{diss},m} \simeq E'_{\text{diss}}/\delta t_{\text{em}}$. We assume that this available energy is distributed among the accelerated protons (ϵ_p), electrons (ϵ_e), and the magnetic field (ϵ_B), with the fractions normalized to $\epsilon_p + \epsilon_e + \epsilon_B = 1$. The energy in each component is then $E'_{p,\gamma,B} = \epsilon_{p,\gamma,B} \cdot E'_{\text{diss}}$. We further make the assumption that the non-thermal electrons lose energy quickly via synchrotron radiation, implying that gamma rays will carry a comparable amount of energy $\epsilon_\gamma \simeq \epsilon_e$. To fix the fractions, we assume a baryonic loading of $\epsilon_p/\epsilon_\gamma = 10$. We further assume equipartition between photons and the magnetic field resulting in the fractions $\epsilon_\gamma = \epsilon_B = 1/12$ and $\epsilon_p = 10/12$.

We do not treat the acceleration of protons explicitly, instead we assume a power-law injection spectrum motivated by Fermi shock acceleration $n'_p(E'_p) \propto (E'_p)^{-2} \exp(-E'_p/E'_{p,\mathrm{max}})$. The maximal energy $E'_{p,\mathrm{max}}$ is determined as the energy where the acceleration rate $t'^{-1}_{\mathrm{acc}} \simeq c/R'_L$ (where R'_L is the Larmor radius) equals the losses due to photomeson interactions and adiabatic expansion. The spectrum is then normalized to the available total energy with the fraction ϵ_p.

We do not compute the photon spectrum explicitly by self-consistent radiation calculations. Instead, we assume a broken power-law motivated by observations

$$n'_\gamma(\varepsilon') \equiv \frac{dn'_\gamma}{d\varepsilon'} = C'_\gamma \begin{cases} (\varepsilon'/\varepsilon'_{\mathrm{break}})^{-\alpha} & \text{, if } \varepsilon' < \varepsilon'_{\mathrm{break}} \\ (\varepsilon'/\varepsilon'_{\mathrm{break}})^{-\beta} & \text{, if } \varepsilon' \geq \varepsilon'_{\mathrm{break}} \end{cases} \tag{6.8}$$

with spectral indices $\alpha = -1$ and $\beta = -2$ below and above the break energy $\varepsilon'_{\mathrm{break}} = 1$ keV, respectively. We assume a maximal photon energy of $\varepsilon'_{\mathrm{max}} = 10^6$ GeV, which is reduced for high luminosities when $\gamma\gamma$-annihilation is efficient. The normalization is performed equivalent to the proton case with the fraction ϵ_γ. While the value of the maximal energy is an ad-hoc choice, the energy densities depend (at most) logarithmically on the maximal and minimal energies.

The coupled differential equation system for this model is mathematically similar to the UHECR propagation problem in Section 3.2, however with a static photon-targed field and additional processes such as particle escape and synchrotron losses. It is evolved to the steady state using the current version of NEUCOSMA [124], which takes photo-hadronic interactions, photo-disintegration, adiabatic cooling and synchrotron losses into account. The neutrons escape freely, whereas the protons are magnetically confined and escape only from the boundaries (within their Larmor radius), referred to as “direct escape” [172]. This assumption implies that close to the photosphere the UHECR emission is dominated by neutron escape, and for large collision radii by direct proton escape [254]. The mechanism for UHECR escape is currently discussed, and it is likely that several competing components contribute; see discussion in Zhang et al. [133]. The implementation for secondary particle emission is described in great detail in Biehl et al. [124]. The results presented in Section 6.4 were first published in Bustamante et al. [1], where we used an older version of the radiation code, in which the energy budget was normalized to the steady state density instead of the injection luminosity. The (more consistent) normalization to the injection luminosity leads to slightly lower UHECR fluxes, see the appendix of Biehl et al. [124] for a detailed comparison. The updated version of NEUCOSMA can also treat the injection of heavier nuclei treating the full nuclear cascade. Therefore, all results from Bustamante et al. [1] shown in this chapter are updated using the current code version and extended by examples for UHECR nuclei injection. The results in Rudolph et al. [4] were already done with the updated code.

Note that our model does not account for the emission from sub-photospheric collisions. We define these as collision for which the optical depth to Thomson scattering τ_T is larger than one. We estimate τ_T by assuming that the number of thermal electrons is equal to the density of thermal baryons $N_e = N_p \simeq m_{\text{shell}}/m_{\text{proton}} = E_{\text{kin,shell}}/(\Gamma m_{\text{proton}})$. With the commoving electron number density $n'_e = N_e/V'_m$ and shell width l'_m, we then get:

$$\tau_T = \sigma_T n'_e l'_m = \sigma_T \, \frac{N_e}{4\pi R^2_{\text{coll}}} \tag{6.9}$$

Note that we use the proton mass here to estimate the electron number density. For nuclei the ratio of nucleons to electrons is modified only by $Z/A \approx 1/2$, which is negligible in most cases. For large values of τ_T the photons scatter inside the shell, producing a thermal spectrum that differs significantly from our assumption of a fixed broken power-law. Even if cosmic-rays were accelerated efficiently in this range, they would not not reach UHECR energies due to the high radiation densities. There might however be a significant contribution of sub-photospheric collisions to the neutrino flux. Our neutrino flux estimate should therefore be considered as a minimal prediction. We chose examples in which the fraction of sub-photospheric collisions is small to reduce the impact of this effect. For the neutrino fluxes, we show a sub-photospheric extrapolation as shaded areas, for which we still use a non-thermal photon spectrum even though it is not fully realistic.

For the time structure of the burst, we assume that each (super-photospheric) collision emits a a gamma-ray pulse. We parameterize the shape of the pulses (in the observers frame) with a fast rise and exponential decay normalized to the energy output, see Bustamante et al. [1] for details. This shapes follows from the curvature of an emitting spherical shell. The observed light curve of the burst is then the superposition of all pulses propagated to Earth.

6.3 Stochastic baseline model

We will explore the impact of the GRB engine model in Section 6.4 and the collision model in Section 6.5. In both cases we compare to a baseline model with a purely stochastic engine with Lorentz factors sampled from a log normal distribution. The relevant features of the multi-collision model will therefore be introduced in detail for the baseline model in this section.

In this baseline model we assume that the initial in Γ distribution is sampled from a log-normal distribution:

$$\ln\left(\frac{\Gamma_k - 1}{\Gamma_0 - 1}\right) = A \cdot x \,, \tag{6.10}$$

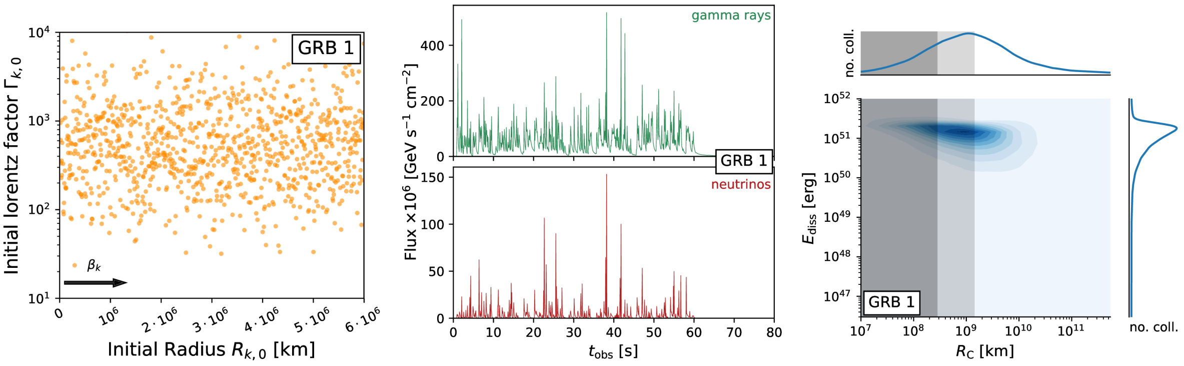

(a) Initial shell setup (b) Lightcurve (c) Dissipation radii

Figure 6.4: Collision dynamics of the stochastic baseline model GRB1 with E_γ, iso $= 5.7\cdot10^{52}$ erg and constant energy injection: (a) Initial Γ distribution of the shells sampled from Eq. (6.10) with $A = 1$, $\Gamma_0 = 500$ and $\delta t = 10^{-2}$s. (b) Synthetic light curve in gamma-rays and neutrinos considering all collisions beyond the photosphere. Each peak is normalized to the luminosity of the collison $E_{\rm coll,diss}/\delta t_{\rm em}$. (c) Distribution of collisions by their dissipated energy over the collision radius. The density profile was created by computing 20 different random seeds for the initial setup. Left and middle panel updated from Bustamante et al. [1].

where x is sampled from a Gaussian $P(x)dx = \exp(-x^2)/\sqrt{2\pi}dx$. The distribution is fully defined by the bulk Lorentz factor Γ_0 and spread A. The spatial separation between shells is assumed to be equal and the same as their width ($L_k = l_k = \delta t_k \cdot c$, with a temporal separation of $\delta t_k = 0.01$s for relativistic shells). We choose the burst parameters as $\Gamma_0 = 500$ and $A = 1$ for 1000 shells with a constant energy outflow of $E_{k,\rm kin} = 10^{52}$ erg. The actual observable quantity is the gamma ray output $E_{\gamma,\rm tot}$. It varies with the efficiency $\varepsilon = E_{\rm diss,tot}/E_{\rm kin,tot}$, which is calculated from the burst evolution. We therefore renormalize $E_{k,\rm kin}$ after the burst simulation to give a fixed gamma-ray output of $E_{\gamma,\rm tot} = 10^{53}$ erg (including both sub- and super-photospheric collisions). For this setup, about 40% of the energy is dissipated in sub-photospheric collision, the observed $E_{\gamma,\rm tot}$ is therefore reduced by this percentage.

The fireball dynamics are shown in Fig. 6.4: The left panel shows the initial shell setup, the middle panel the resulting light curve in gamma rays (green) and neutrinos (red) and the right panel the distribution of collisions by their dissipated energy over radius. The light curves in both particle species are dominated by fast time variability and show no other prominent features. However, there is also no direct correlation between the height of peaks in either light curve. The dissipation radius shows a broad distribution around an average radius. The dissipated energy per collision is relatively peaked around a value of $E_{\rm diss} = 2 \cdot 10^{51}$ erg, with

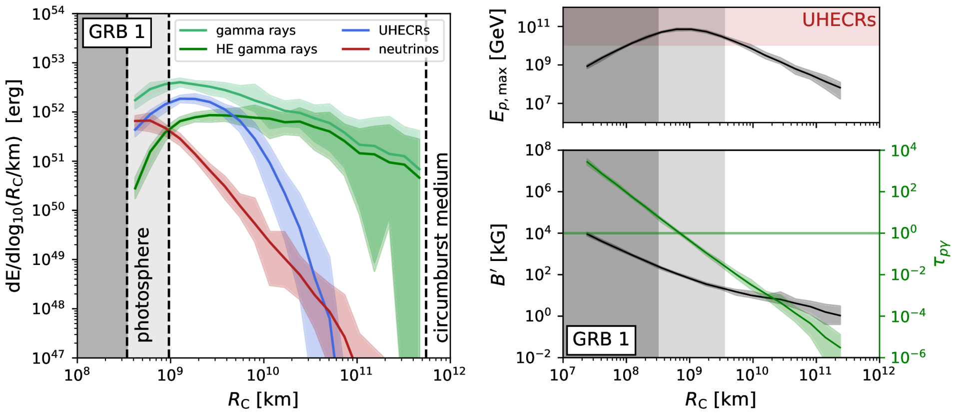

(a) Particle prodution regions

(b) Maximal UHECR energy

Figure 6.5: Radiation output for the stochastic baseline model GRB1: (a) Energy dissipated in neutrinos, gamma rays and UHECRs as a function of radius. For UHECRs we consider cosmic rays above 10^{10} GeV in the source frame. The photosphere is marked by grey shaded areas. In the dark-shaded areas, all collision are sub-photospheric, while in the light shaded area both sub- and super-photospheric collisions occur. Sub-photospheric collision are not considered for particle radiation. (b) The average maximal UHECR energy (upper panel) as well as the average magnetic field and optical depth to $p\gamma$-interactions (lower panel, different y-axes marked by colors) all as a function of radius. The shaded regions in all panels show the fluctuations for 20 representations of the initial stochastic shell setup.

later collisions dissipating less energy on average. Note that the whole fireball moves with a bulk speed of $c \cdot \beta \approx c$ towards the observer, while the light still has to reach the observer with c. There is therefore almost no correlation between the radius of the collision and the time of its observation. Instead, the time structure of the light curve depends on the relative position of the collision inside the fireball. Due to the random structure of the initial setup, neutrino-bright collisions, that happen at small radii, are therefore equally distributed in the light curve. This will change for more structured initial setups in Section 6.4.

The dependence of the radiated particles on the radius is shown in Fig. 6.5. For this plot we compute 20 different representations drawn from the initial distribution to account for the stochastic fluctuations in the initial sample. In these figures, we show the average as a solid curve and the region between the edge cases as shaded bands.

The left panel (a) shows the differential energy dissipation in the engine frame for the different messengers: neutrinos, UHECRs (defined here above 10^{10} GeV) and gamma rays (with high-energy gamma rays above 1 GeV). In the dark-shaded region, all collisions occur below the photosphere, while in the light-shaded region, the collision may be above or below the photosphere (depending on the other collision parameters). Neutrinos originate near the photosphere due to the high radiation density; UHECRs prefer intermediate collision radii since high magnetic fields are required for efficient acceleration – but not as high that radiation losses (such photo-hadronic interactions) limit the maximal energy; gamma rays trace the region where most energy is dissipated, i.e. they correspond to the distribution of collisions in the right panel of Fig. 6.4. The high-energy gamma rays escape only at larger radii, where their maximal energy is not limited due to pair-production.

Since the astrophysical messengers originate dominantly from different radii, the initial setup of shells and the collision model have important impact because they generate the distribution of collisions over radii. The statistical fluctuations from the initial shell setup (shaded areas) imply some variability in this picture, however they are small in the region where most energy is dissipated and do not change the qualitative picture. Modifications to the model assumptions will have a stronger qualitative impact.

The evolution of the maximal UHECR energy is shown in Fig. 6.5 (b). The lower panel shows the magnetic field B' (left axis, black curve) and the optical thickness to photo-hadronic interactions $\tau_{p\gamma}$ evaluated at $E_{\mathrm{p,max}}(R_{\mathrm{coll}})$ (right axis, green curve). Both scale with the collision radius, $\tau_{p\gamma} \sim R_{\mathrm{coll}}^{-2}$ and $B' \sim R_{\mathrm{coll}}^{-1}$. This scaling explains the shape of the maximal UHECR energy in the upper panel. At low radii, the maximal energy is limited by interactions, while at high radii acceleration is not efficient as the low magnetic fields cannot confine the particles. The highest UHECR energies are therefore reached at intermediate radii. As we only consider UHECRs above 10^{10} GeV (red shaded range), the total UHECR output in panel (a) has a similar shape as the maximal energy.

6.4 Implications of the GRB engine properties

To illustrate the impact of the GRB engine, we now choose different examples of the initial Γ setup. A purely random distribution in Γ always leads to light curves dominated by fast time variability as energetic collisions happen mostly between neighboring shells. We therefore now discuss more structured setups for the initial engine. In Bustamante et al. [1] we show a total of ten different cases (including the appendix). We limit the discussion here to the three cases that show the most qualitative differences, GRBs 3-5. We stick to the numbering from Bustamante et al. [1] for easier comparison. Note that the calculations have been redone with an updated version of the NeuCosmA code, which affects the particle output slightly. The most relevant change is a corrected factor of $(1+z)$ for the neutrinos leading to 3 times lower integral energies and light curves, while the observed fluxes not affected.

Table 6.1: Description of the initial setup of the GRBs 1 and 3–5

Model	$\Gamma_{0,1}$	$A_{\Gamma,1}$	$\Gamma_{0,2}$	$A_{\Gamma,2}$	T_p	$E^{\text{iso}}_{\gamma,\text{norm}}$	Description
1	500	1.0	–	–	–	10^{53} erg	Fixed Γ and A_Γ; baseline model
3	50	0.1	500	0.1	0.34	10^{53} erg	Sawtooth Γ (linear slowdown three times) with narrow distribution
4	50	0.1	500	1.0	0.2	10^{53} erg	Oscillating Γ (five periods) with increasing distribution width
5	50	0.1	500	0.1	0.2	10^{53} erg	Oscillating Γ (four periods) with lower amplitude increasing and narrow distribution

Common values for all models: $N_{\text{sh}} = 1000$, $\delta t_{\text{eng}} = 10^{-2}$ s, $d = l = c \cdot \delta t_{\text{eng}}$, $r_{\text{min}} = 10^3$ km, $r_{\text{dec}} = 5.5 \cdot 10^{11}$ km, $z = 2$, $\epsilon_e = \epsilon_B = 1/12$, $\epsilon_p = 5/6$, $\eta = 1.0$ (acceleration efficiency [172]). The period T_p for the oscillating cases refers to Γ changing between $\Gamma_{0,1}$ and $\Gamma_{0,2}$, and A_Γ changing between $A_{\Gamma,1}$ and $A_{\Gamma,2}$; T_p is a fraction of the total number of emitted shells. This means that Γ and A_Γ change between first and second value with a factor $\sin^2(k/(N_{\text{sh}} \cdot T_p) \cdot \pi)$, where k is the index of shell ($1 \leq k \leq N_{\text{sh}}$). For GRB 5, the lower amplitude increases from $\Gamma_{0,1}$ to $\Gamma_{0,2}$ linearly with k.

Table 6.2: Output parameters of the simulated GRBs 1 and 3–5.

Model	N_{coll}	t_{v} [ms]	T_{90} [s]	ε [%]
GRB 1	988 ± 2	55 ± 0	53.9 ± 0.5	25.8 ± 0.8
GRB 3	949 ± 4	35 ± 0	33.0 ± 0.4	10.0 ± 0.2
GRB 4	983 ± 3	53 ± 0	51.7 ± 0.3	21.6 ± 0.3
GRB 5	990 ± 2	59 ± 1	58.2 ± 0.6	10.9 ± 0.2
	$E^{\text{iso}}_{\gamma,\text{tot}}$ [10^{52} erg]	$E^{\text{iso}}_{p,\text{tot}}$ [10^{52} erg]	$E^{\text{iso}}_{\nu,\text{tot}}$ [10^{51} erg]	$\frac{E^{\text{iso}}_{\nu,\text{tot}}}{E^{\text{iso}}_{\gamma,\text{tot}}}$ [%]
GRB 1	5.92 ± 0.29	1.89 ± 0.14	4.82 ± 0.35	8.2 ± 0.6
GRB 3	6.34 ± 0.18	0.86 ± 0.09	1.17 ± 0.25	1.8 ± 0.4
GRB 4	5.31 ± 0.23	1.08 ± 0.06	3.11 ± 0.17	5.9 ± 0.3
GRB 5	8.94 ± 0.03	0.43 ± 0.02	0.16 ± 0.01	0.2 ± 0.0

The parameters are: variability timescale (t_{v}), duration T_{90} (defined as the time where 90% of the energy is dissipated), total energy emitted as gamma rays ($E^{\text{iso}}_{\gamma,\text{tot}}$), as protons ($E^{\text{iso}}_{p,\text{tot}}$), as neutrinos ($E^{\text{iso}}_{\nu,\text{tot}}$), ratio between neutrino and gamma-ray energies ($E^{\text{iso}}_{\nu,\text{tot}}/E^{\text{iso}}_{\gamma,\text{tot}}$), and overall emission efficiency ε. Energies are computed using super-photospheric collisions only. Only protons in the UHECR energy range above $> 10^{10}$ GeV are counted. The efficiency ε is defined as the ratio of total energy dissipated by all types of particles (gamma rays, protons, neutrinos) to the total kinetic energy initially available.

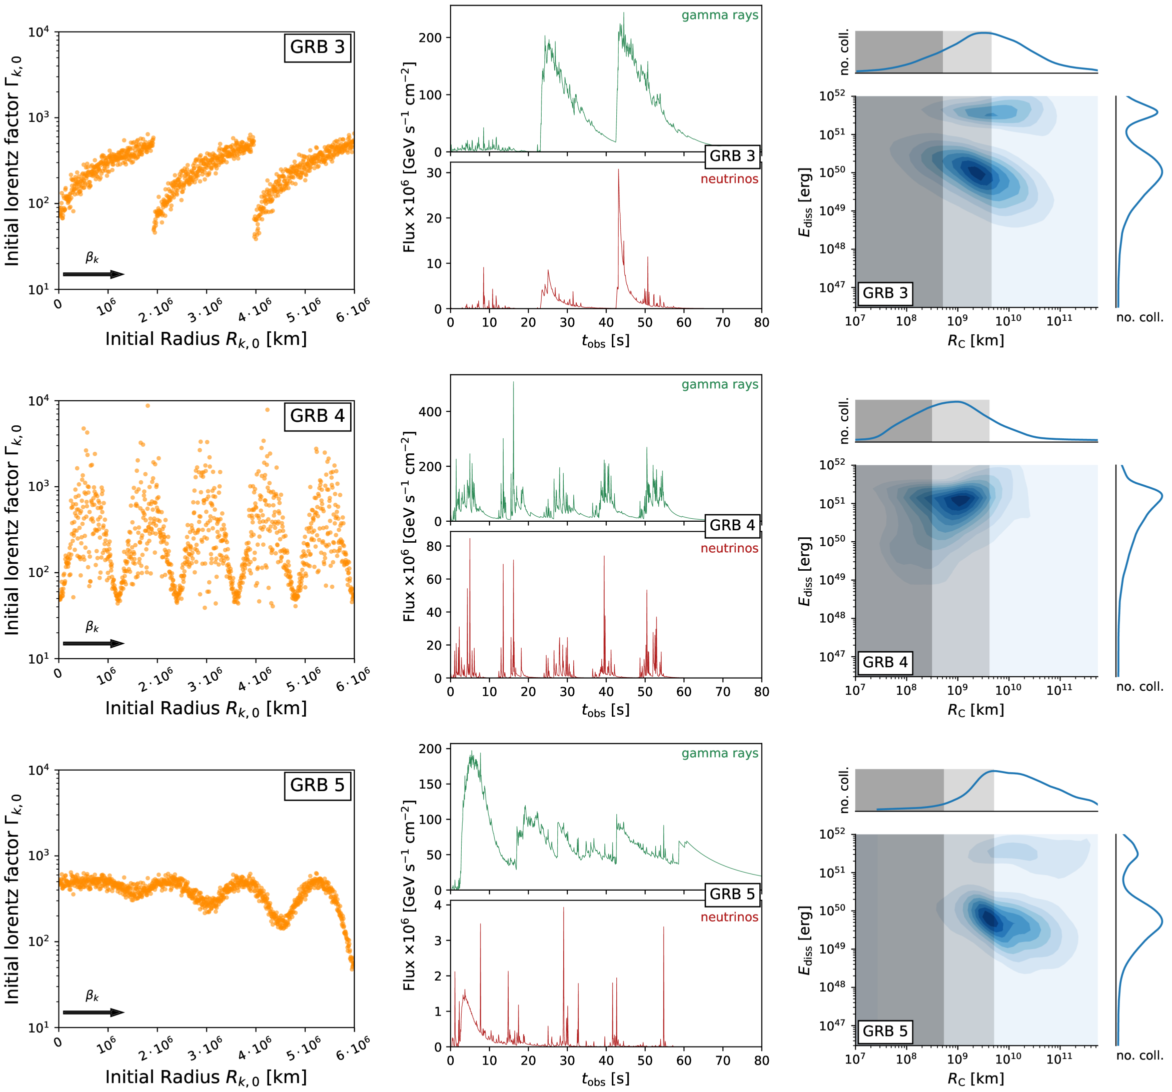

Figure 6.6: Initial shell setup (a), resulting light curves (b) and dissipated energy over radius (c) for the GRBs 3–5 with the setup generated by the parameters in Tab. 6.1. The format is the same as shown for the baseline model in Fig. 6.4. Left and middle row updated from Bustamante et al. [1].

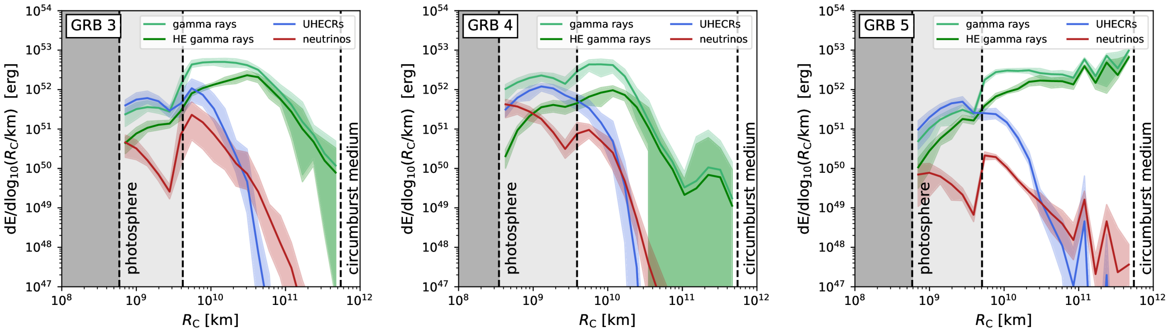

Figure 6.7: Particle production regions for neutrinos, UHECRs and gamma rays for the GRBs 3–5. The format is the same as shown for the baseline model in Fig. 6.5, left panel.

The initial shell setup and resulting light curve and dissipation radii are shown in Fig. 6.6. The parameters used to generate the initial shell setups are shown in Tab. 6.1 Note that there is not always a simple correlation between the shell setup and the light curve. Fast regions with low variability need slower shells in front to collide with, otherwise they will dissipate only little energy and not affect the observation significantly. Fast regions with strong variability will however dissipate energy fast, which slows them down. They can then collide with other regions later in the fireball evolution.

GRB 3 consists of three pulses that each have a slowdown in Γ. This leads to three separated pulses in the light curve. The overlaid fluctuations in Γ are smaller, leading to less a pronounced time-variability than in the baseline model. Due to the slowdown of the first pulse, the first emitted shells do not have a target to collide with, leading to low dissipation efficiency. The two later pulses can instead collide with the tail of the preceding pulse. The first pulse in the light curve thus has a much smaller amplitude than the two later pulses. GRB 4 has an oscillating period, where five regions with high amplitude and fluctuations alternate with slower regions. This leads to five distinct peaks in the light curve, each overlaid with strong time-variability. GRB 5 also shows a periodic amplitude, where however the lower amplitude increases during emission. The overlaid fluctuations are small. This leads to a more structured light curve, which is dominated by a large initial peak, produced when the faster shells in the rear run into the front.

The structure in the initial setup can be seen in the dissipation radii in Fig. 6.6, right column. GRBs 3 and 5 both have a bimodal distribution: A large number of collisions have low E_{diss}, they are due to the overlaid fluctuations. Most of the energy is dissipated in a few collisions, which happen when the larger structures in the outflow collide. The later collisions happen at slightly larger radii. As GRB 4 has structured pulse overlaid with strong variability, it has a broad distribution in the collision parameters, which is a mix of GRB 1 and 3.

As discussed in Section 6.3, the different messengers come from different radii, as the production efficiency scales differently with radius for neutrinos and UHECRs. While this relation qualitatively holds in the more structured cases, the effect is more complicated due to the bimodal distribution of collisions. The particle production regions for GRBs 3–5 are shown in Fig. 6.7. Close to the photosphere the particle output looks similar to GRB1 for all cases, i.e. $R_{\mathrm{coll}} \approx 5 \cdot 10^8 - 5 \cdot 10^9$ km the neutrino flux peaks directly at the photosphere and falls steeply while UHECRs at slightly higher radii peak around $R_{\mathrm{coll}} = 10^9$ km. This initial dissipation is due the fluctuations in Γ, which are smaller for GRB 3 and 5 and therefore have a lower energy output at these radii. For GRBs 3 and 5, most of the energy is instead dissipated around 10^{10} km, where the different regions of the fireball collide. Most neutrinos are produced in these more luminous collisions, which are at high enough radii to also accelerate UHECRs. The separation between production radii is therefore not as pronounced in structured GRBs. For GRB5 most of the energy is actually dissipated at large radii, such that the overall neutrino output is small. For GRB 4 there is a second bump as well, but it is not as pronounced, as a lot of energy is already dissipated by the early collisions. We will discuss this further for the expected neutrino fluxes in Section 6.4.1 and the UHECR output spectrum in Section 6.4.2.

The non-linear evolution of the fireball leads to an ambiguity in the interpretations of the light curve in terms of initial shell distribution. We chose here distinct examples to discuss the qualitative effects. With enough fine-tuning one could in principle find multiple, qualitatively-different initial setups that produce similar light curves. There is no straightforward way to infer the engine properties from the light curve.

6.4.1 Impact on the expected neutrino fluxes

The IceCube collaboration has derived a stacking limit on the contributions of GRBs to the diffuse neutrino flux [24, 25]. To compare to this limit, we scale the neutrino fluence prediction for a single GRB (F_{ν_μ}) from our calculations by a rate of $\dot{N} = 667\ \mathrm{yr}^{-1}$, i.e. $J_{\nu_\mu} = \mathrm{F}_{\nu_\mu} \cdot \dot{N} \cdot (4\pi)^{-1}$, the same rate that was assumed in the IceCube analysis. As this assumes all GRBs to be identical, it only gives a rough estimate of the exclusion limit. A detailed prediction of the neutrino flux would have to model each GRB in the observed catalogue individually requiring some kind of automatic algorithm, which is difficult due to the large ambiguity in the model parameters. Still, the assumption of equal GRBs gives an estimate of whether specific types of GRBs can be excluded as the dominating contribution.

Fig. 6.8 shows the neutrino fluxes calculated for the baseline model (GRB1) and the alternative engine setups GRBs 3-5. The total energy output in different species is given in Tab. 6.2. Note

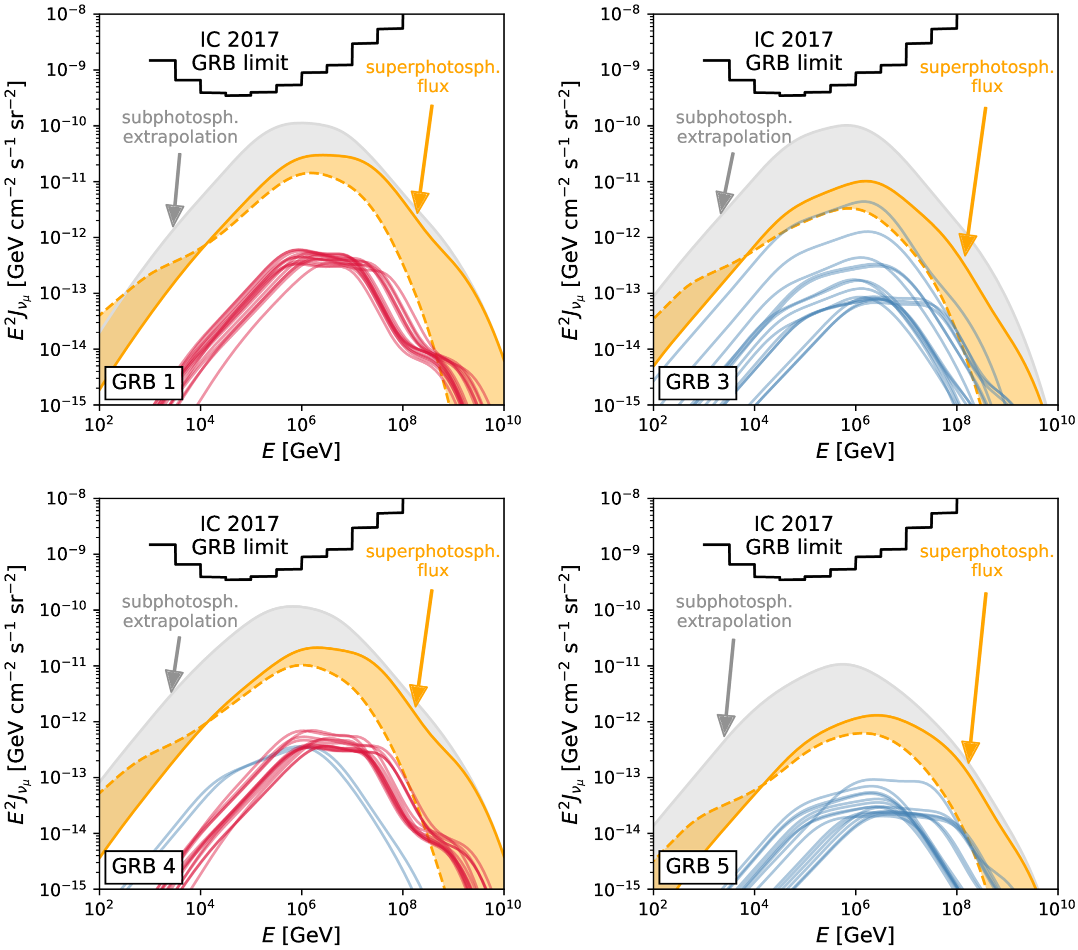

Figure 6.8: All-sky quasi-diffuse $\nu_\mu + \bar{\nu}_\mu$ fluxes for the simulated multi-zone GRBs 1 and 3–6. The flux prediction for pure proton injection (orange, solid) and pure iron injection (orange, dashed) only considers super-photospheric collisions. A realistic mixed composition would produce a flux in the range indicated by the orange shaded region. The grey shaded region gives the potential contribution from sub-photospheric collisions. The dominant contributions from individual collisions are shown as thin curves, corresponding to cases where the optical depth to photo-hadronic interactions $\tau_{p\gamma}$ is larger (red) or smaller (blue) than unity. The IceCube 2017 upper limit is taken from the latest GRB stacking search using tracks from the Northern and Southern Hemisphere [25]. Figures updated from Bustamante et al. [1]

that the initial kinetic energy for each GRB was renormalized to result in comparable gamma-ray outputs of $\sim 5 \cdot 10^{52}\,\mathrm{erg}$. The level of the neutrino fluxes are therefore impacted only by the different engine setups. Slight differences to Bustamante et al. [1] in these plots come from the new version of the NeuCosmA code used. A line for the neutrino production for pure iron injection is added (dashed line).

We show the prediction for the neutrino flux for super-photospheric collisions (orange curve). As mentioned before our radiation model is not accurate for sub-photospheric collisions, where the photon spectrum thermalizes and UHECR are not accelerated as efficiently. We still show an extrapolation of our model for these collisions up to an optical depth of $\tau_{p\gamma} < 20$ (grey curve), to get an estimate of the maximal contribution. A realistic sub-photospheric model would likely produce less neutrinos than this extrapolation due to the lower peak photon energy of a thermal spectrum. We also show the individual collisions that contribute most to the total flux. Their color indicates whether they are optically thin (blue) or thick (red). The two fluxes for the two GRBs that are dominated by fast variability (GRBs 1 and 4) produce neutrino fluxes on similar levels of $(2-3) \cdot 10^{-11} \frac{\mathrm{GeV}}{\mathrm{cm\,s\,sr}}$. The flux here is dominated by optically-thick collisions. The fluxes for the other two GRBs 3 and 5, are lower due to their more disciplined engine. The smaller variability in Γ leads to higher collision radii and therefore optically-thin collisions. For GRB 3 there are only a few dominant collisions that lead to a flux level of $1 \cdot 10^{-11} \frac{\mathrm{GeV}}{\mathrm{cm\,s\,sr}}$, which is close to the level of the baseline model GRB1. These happen at the onset of each pulse, as can be seen from the neutrino light curve in Fig. 6.6. The engine of GRB5 is more disciplined, leading to the lowest peak flux level around $1 \cdot 10^{-12} \frac{\mathrm{GeV}}{\mathrm{cm\,s\,sr}}$. For iron injection the neutrino flux is lower by a factor of ~ 2 and peaks at slightly lower energy for all models. A realistic mixed composition would produce a neutrino flux in between the proton and iron result as indicated by the orange shaded area.

Note that most observed long GRB light curves do not have the pronounced features seen for GRB 3 and 5 and are instead dominated by fast variability. The expected neutrino flux level for the whole population is therefore likely to be relatively robust at the level predicted for GRB 1 and 4. Still these predictions are significantly lower than what would be expected from one-zone models, where only a single collision with an intermediate radius is calculated and then upscaled to the total emission. The average parameters used in this calculations are derived from gamma-ray observations, which are not necessarily representative for neutrino calculations. While one-zone models are already strongly challenged by the current neutrino limits, the more realistic multi-zone limits relax this limit.

Note that we did not fit the resulting ejected UHECRs to the observed spectrum for this study but instead chose a baryonic loading of 10. The energy budget in UHECRs (see Tab. 6.2), however, suggests that this value has to be increased by a factor of a few to reach required total energy output per GRB of $\sim 3 \cdot 10^{53}$ ergs [224]. The neutrino flux will therefore also be increased by the same factor. The qualitative impact on the shape of the UHECR output will be discussed in the following section.

6.4.2 Impact on the UHECR spectrum

The shape and normalization of the UHECR spectrum depends not only on the radiation model, but also on the distribution of the collision parameters and therefore by extension on the initial shell setup. We therefore discuss the UHECR output of the GRBs 1 and 3–5. While the main results are for pure proton composition, they still give some qualitative insight on the parameters required to power UHECRs. We also show the spectra for pure iron injection to illustrate the other extreme case. A realistic fit to the Auger composition would require a mixed injection.

The energy output for the models GRB 1 and 3–5 in Tab. 6.2 ranges from $0.5 - 2 \cdot 10^{52}$ erg, this means they are all are relatively efficient UHECR emitters. The impact of the engine setup on the total UHECR output is less severe compared to the neutrino output. To power UHECRs an energy output of $3 - 11 \cdot 10^{53}$erg is needed depending on the exact source evolution [224]. This implies an increase in baryonic loading of a factor ~ 10, where the simulations presented here were done with a fixed baryonic loading of $\epsilon_p/\epsilon_\gamma = 10$. This would increase the cosmogenic neutrino fluxes by roughly the same factor, with some smaller non-linear effects due to e.g. a shift in the photosphere. Still, these would be slightly below the current IceCube GRB limit. Considering only the rough energy budget, it is therefore still feasible that GRBs can be the sources of UHECRs.

The GRB models also produce a spectral shape and composition that can account for the UHECR observations after propagation. As seen by the parameter scans in Chapter 5, these should be hard spectra with a maximal rigidity of $\sim 2 \cdot 10^9$ GV for GRB evolution (roughly SFR). The spectra for our GRB simulations with pure-proton injection are shown in Fig. 6.9, separated by their escape mechanism: In the 'direct escape' (red) case, a fraction of protons escapes directly through free streaming. This escape is limited by their Larmor radius, which scales with energy, leading to hard escape spectra. For 'neutron escape' (marked blue) photo-hadronic interactions of protons produce neutrons, which are not magnetically confined and escaped freely. They then decay back to protons outside of the source.

Neutron escape dominates at high luminosities, where also neutrino production is efficient. As seen in Fig. 6.9 the escape spectrum is in all cases dominated by direct escape. The relative contribution of neutron escape depends on the model in a similar way to the amplitude of neutrino fluxes, since for collisions with high luminosity, both neutron escape and neutrino production are effective. For GBRs 1 and 4 the strong fluctuations in Γ lead to smaller average

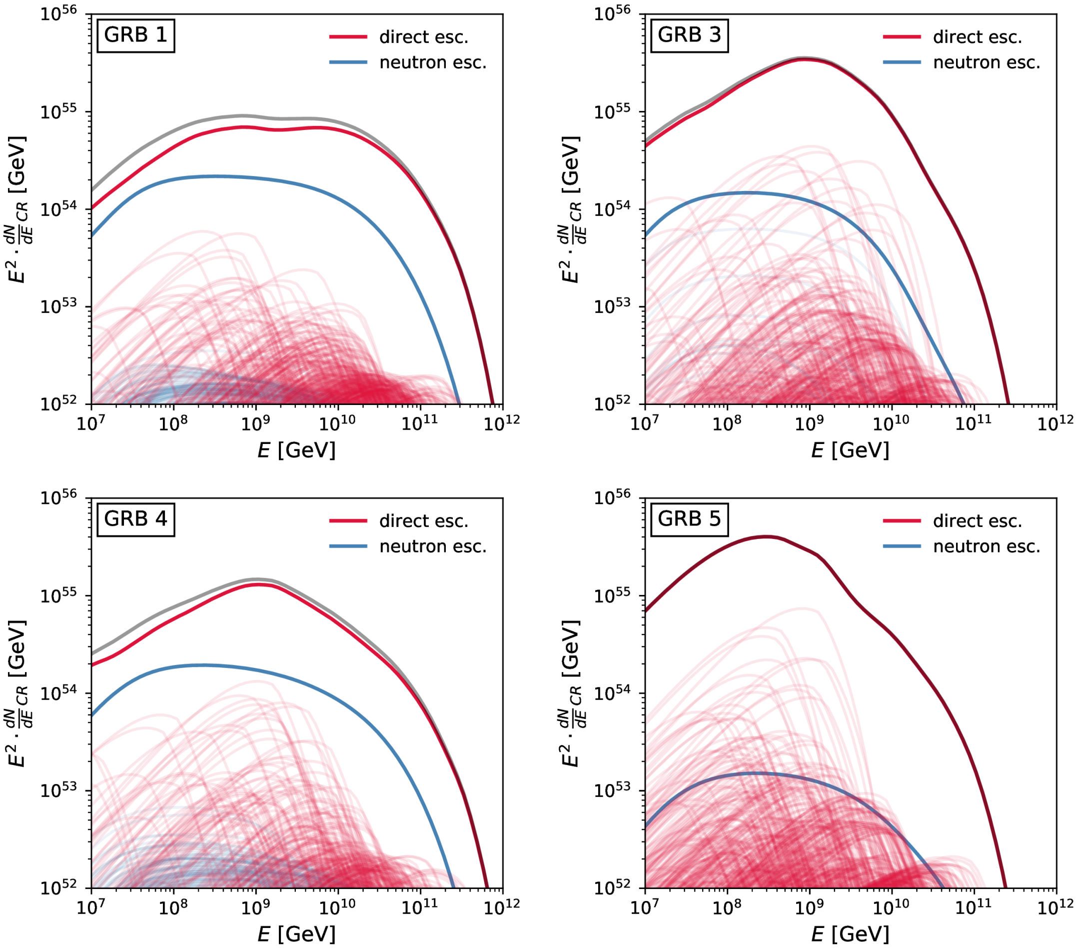

Figure 6.9: Ejected cosmic ray spectrum for GRBs 1 and 3–5 for pure protons composition. The spectrum is shown for direct escaping protons (red) and neutron escape (blue). The total spectrum (grey) is in all cases dominated by direct escape. The spectra of each individual collision are shown as thin curves with the same color scheme.

collision radii. This causes a significant contribution of neutron escape spectra at a level of 1/5 of the total flux. Still, the overall shape is only slightly affected. For the other two models with more disciplined engines, the contribution from neutron escape is negligible.

The maximal energy for UHECRs scales with the collision radius. At low radii, it is limited by interactions due to high luminosities, while at high radii the acceleration rate decreases due to decreasing magnetic fields. There is therefore a sweet spot for reaching the highest maximal energies at intermediate radii. The total spectrum is however a superposition of all collisions. Compared to simple one-zone models, which would have a shape similar to the spectra of a single individual collision, the shape of the UHECR spectrum is therefore significantly altered

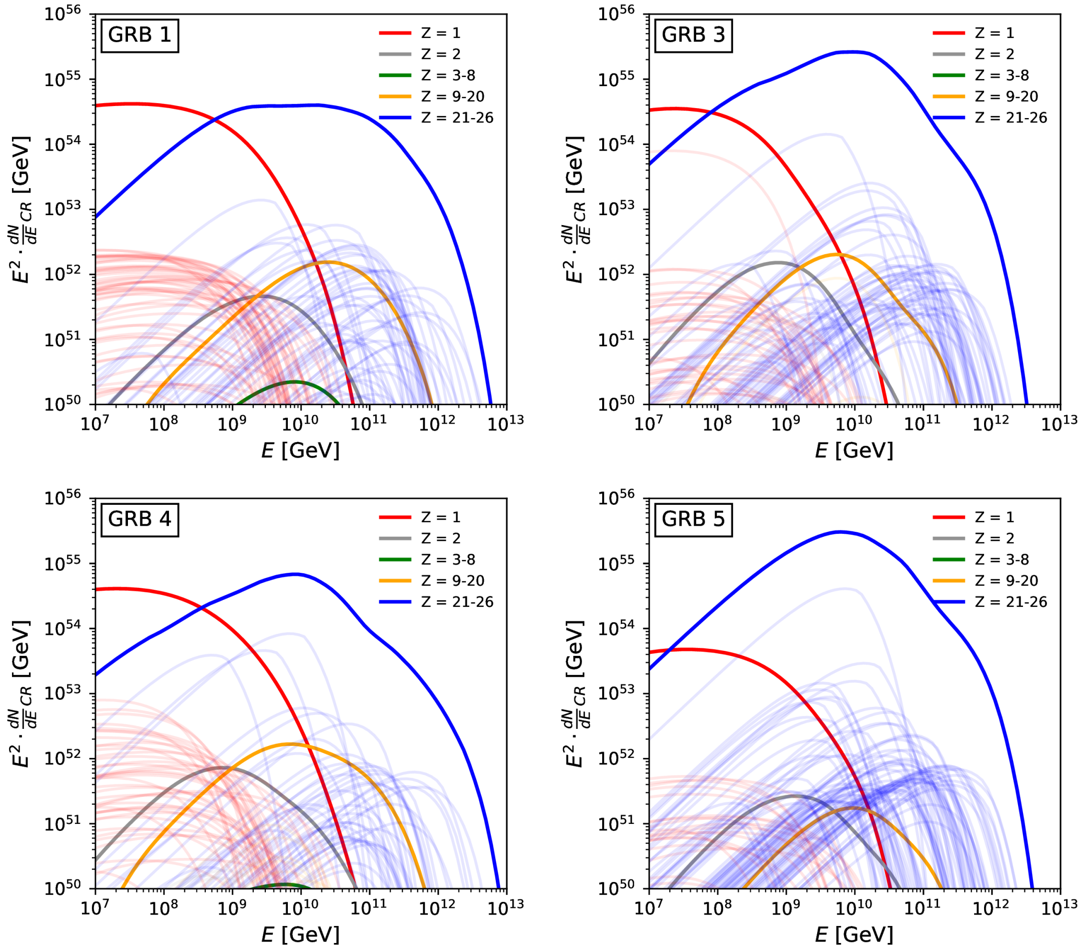

Figure 6.10: Ejected cosmic ray spectrum for GRBs 1 and 3 for pure iron composition. The spectrum is grouped by the particle charge (see legend). The spectra of each individual collision are shown as thin curves with the same color scheme.

in our multi-zone model.

As seen in Fig. 6.9 for direct escape there are a few collisions with high luminosity, but low maximal energy and many collisions with lower luminosity, but larger maximal energy. In the purely stochastic case, their contributions balance each other producing a total spectrum that is significantly softer than the individual spectra. The maximal energy of $> 10^{10}$ GeV reached in this case is actually higher than the maximal rigidity required to fit UHECRs, while the overall spectrum is too soft. The more disciplined engines produce collisions at higher radii. Their total spectra therefore peak at lower energies. The collisions with smaller radii but higher maximal energy contribute less, leading to a smooth cutoff of the total spectrum.

Of course, a realistic model accounting for the Auger composition would require nuclei injected into the source. As the other extreme case, we show the spectra for pure iron injection in Fig. 6.10. This reveals that the qualitative effect for nuclei injection are similar to those discussed for protons. At low radii, the maximal energies are limited by disintegration, which produces a nuclear cascade resulting in a large flux of secondary protons. The relative amplitude of this proton component is higher for the GRBs with stronger variability. The spectrum of escaping iron is dominated by optically thin collisions, where the maximal energy is limited by the acceleration rate. The escape spectra of intermediate masses are small. A fit to the observed composition would therefore require a mixed mass injection. The acceleration and escape rates are determined by the Larmor radius, such that the maximal energy scales with charge if losses are sub-dominant (Peters cycle). Iron injection therefore reaches ~ 26 times higher energies than protons. While the maximal energy is sufficient to fit the observed spectrum of UHECRs, the superposition of the individual spectra leads to a total spectrum that is likely too soft considering the hard spectral indices that are required by the UHECR fit in Chapter 5.

From this qualitative discussion, a more disciplined engine would be required to produce UHECR spectra that could fit the propagated UHECR flux, especially the hard spectral indices. The distribution of collision radii should be relatively peaked, as otherwise the superposition of the different spectra makes the total spectrum too soft as seen in Fig. 6.9. While the GRB examples presented here produce spectra that are likely too soft to fit the observations, it has been demonstrated in Globus et al. [122] that GRB multi-collision models can in principle produce a reasonably good fit to the Auger spectrum and composition. While Globus et al. [122] showed this only for a single disciplined engine setup without variability, there are certainly complex correlations between GRB engine parameters and the required injection composition. At the same time, we demonstrated that the neutrino flux depends on the engine setup, especially its variability. The predicted level of neutrino flux is therefore likely to distinguish between different engine setups.

6.5 Implications of the collision model

As discussed before, one problem of the internal shock model is the low efficiency of the conversion from kinetic to dissipated energy. This is because the kinetic energy that is not dissipated will affect the afterglow, which is constrained by observations. Also, higher kinetic energies require more extreme assumption for the acceleration models, requiring them to either accelerate more bulk mass or to accelerate the same mass to higher energies. Kobayashi and Sari [284] therefore proposed an *Ultra-efficient* model, where in each collision only a fraction of the internal energy is radiated. The remaining energy results in a splitting of the merged shell, with two shells in the final state. While individual collisions are less efficient in this model, the overall efficiency of the fireball is higher, because the two remaining shells can produce more subsequent collisions. In most cases, the fireball will almost fully thermalize, converting almost all free kinetic energy to radiation.

To explore the effect on the observed neutrino fluence, we therefore discussed different models of the individual collisions in Rudolph et al. [4]. In this section, we focus on the *Ultra-efficient model*, where only a fraction of the internal energy in Eq. (6.3) is radiated and the rest leads to a splitting of the shell, i.e. $E_{\mathrm{diss}} = \eta \cdot E_{\mathrm{int}}$ with $\eta = 0.5$. To compare to the baseline model, we also discuss the *Reduced Efficiency* model, in which also only $\eta = 0.5$ of the internal energy is radiated but the remaining internal energy leads to an adiabatic expansion of the merged shell instead.

In Rudolph et al. [4], we also discuss the validity of these two assumptions by performing hydrodynamic simulations using the Pluto code [287]. As discussed before in Kino et al. [288], a splitting into two shells is dependent on the parameters of the two shells. It mainly happens for similar shell masses $m_r \approx m_s$, but large difference in speed $\Gamma_r \gg \Gamma_s$. This is viable for the initial collisions, if one chooses an equal mass setup for the initial shells. To emphasize the model differences, we therefore choose an equal mass setup $m_{\mathrm{init},k} = \mathrm{const.}$ in this section (in contrast to the equal energy setup used in the last section). However, the *Ultra-efficient model* relies on this also happening in later stages of the fireball, where the spread in Γ is already reduced due to previous collisions. We explored a model, in which each collision in the fireball is directly calculated using Pluto. This model will not be discussed further here, both to keep the discussion short and because I did not directly contribute to the Pluto simulations. Note however that the results of this model are qualitative comparable to the *Reduced Efficiency* model, as the *Ultra-efficient* model applies only to a small fraction of the collisions.

For the same initial setup, the *Ultra-efficient* model would produce many more collisions and therefore different observables than in the baseline model. We therefore choose different initial setups for the models to produce comparable observables like burst durations, variability times derived from the light curve, and total gamma-ray luminosities. By construction, all GRBs are normalized to release the same amount of energy in photons in the optically thin regime of $\sim 5 \cdot 10^{52}$ erg. The burst duration ist determined by the initial size of the system (the sum of all initial shell widths and separations), which matches among the different models. The time variability only depends on the number of collisions for a constant burst duration, which scales with the number of initial shells. In the *Ultra Efficient* model, shells do not merge when colliding, leading to substantially more collisions for the same number of initial shells. We compensate for this by reducing $N^{\mathrm{shells}}_{\mathrm{inital}} = 1000$ to 125. As this affects the fraction of subphotospheric collision, we now normalize only the superphotospheric collision to a value of $5.2 \cdot 10^{52}$ erg for comparable

Table 6.3: Qualitative comparison of the different collision models. Here $N^{\text{shells}}_{\text{inital}}$ is the number of shells ejected by the source, $l^{\text{shells}}_{\text{initial}}$ are the initial shell widths (assumed to be equal for all shells) and $N^{\text{shells}}_{\text{post-coll.}}$ is the number of shells produced by a single collision. In each collision, the dissipated energy is given by $\eta \equiv E_{\text{diss}}/E_{\text{int}}$, and the width of the shell(s) after the collision is given by $l^{\text{shells}}_{\text{post-coll.}}$.

	Initial setup		Single collision result		
Model name	$N^{\text{shells}}_{\text{inital}}$	$\frac{l^{\text{shells}}_{\text{initial}}}{c}$	$N^{\text{shells}}_{\text{post-coll.}}$	η	$l^{\text{shells}}_{\text{post-coll.}}$
Reference	1000	0.01s	1	1	$< l_r + l_s$
Reduced Efficiency	1000	0.01s	1	0.5	$= l_r + l_s$
Ultra Efficient	125	0.08s	2	0.5	$l_r,\, l_s$

Table 6.4: Output parameters for different collision models

	ϵ [%]	t_ν [ms]	N_{coll}
Reference	35.8 ± 1.4	55.2 ± 1.3	970.1 ± 3.3
Reduced Efficiency	17.9 ± 0.7	54.8 ± 1.3	976.0 ± 3.3
Ultra Efficient	36.0 ± 4.3	47.5 ± 10.7	1107 ± 220

	$E^{\text{tot}}_{\text{kin}}$ [10^{54} erg]	$E^{\text{iso}}_{p,\text{tot}}/E^{\text{iso}}_{\gamma,\text{tot}}$	$E^{\text{iso}}_{\nu,\text{tot}}/E^{\text{iso}}_{\gamma,\text{tot}}$
Reference	1.75 ± 0.07	0.42 ± 0.03	0.29 ± 0.05
Reduced Efficiency	3.50 ± 0.13	0.56 ± 0.04	0.24 ± 0.05
Ultra Efficient	1.76 ± 0.22	0.62 ± 0.06	0.14 ± 0.06

results. For the baseline model this results in roughly the same norm as normalizing all collision to 10^{53} erg as in the previous section. In all cases, we assume an engine with a constant mass outflow and constant up- and downtimes, thus equal initial shell widths and separations. An overview of the model parameters is given in Tab. 6.3. The results for the different models are shown in Fig. 6.11, values of the resulting output parameters are given in Tab. 6.4.

As the change to an equal-mass setup has a significant quantitative impact, we show the baseline model again for equal-mass in Fig. 6.11a. For two shell collisions, the conversion efficiency per collision $\varepsilon_{\text{coll}}$ is highest for equal masses, where it asymptotically approaches 1. In contrast, the maximal efficiency in the equal energy case is limited to $\varepsilon_{\text{coll}} \leq 0.3$ [288]. For this setup a few of the initial collisions therefore have very high E_{diss}, while the bulk of collision happens at higher radii and lower energy (middle panel). As more energy is dissipated close to the photosphere, the neutrino flux is enhanced relative to the equal energy case. Fluctuations of the initial setup have a stronger impact, as the dissipation is dominated by fewer collisions. Otherwise the result is qualitatively similar to the equal-energy case in Fig. 6.4 and Fig. 6.5.

The results for the *Reduced efficiency* model are shown in Fig. 6.11b, with the baseline model overlaid as dashed curves. Deviations are visible at large radii, where the magnetic field is lower,

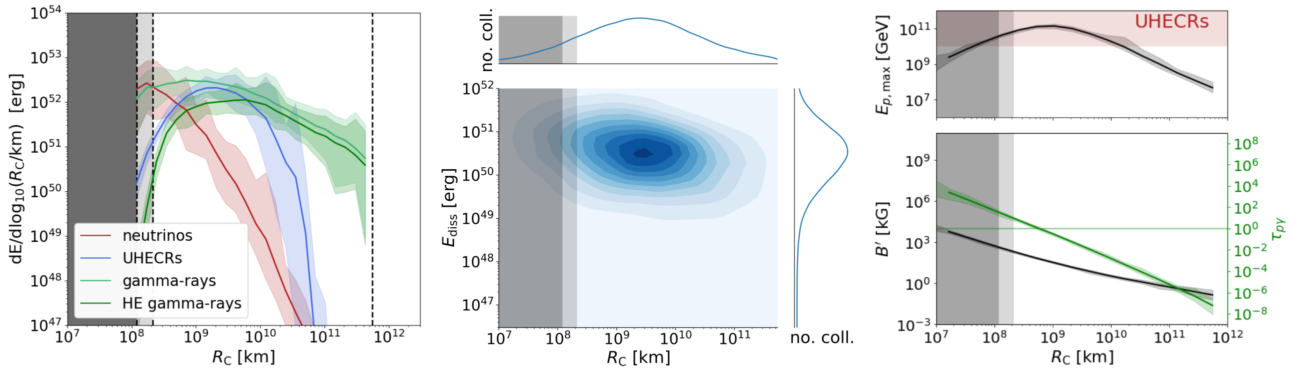

(a) Baseline model with equal mass

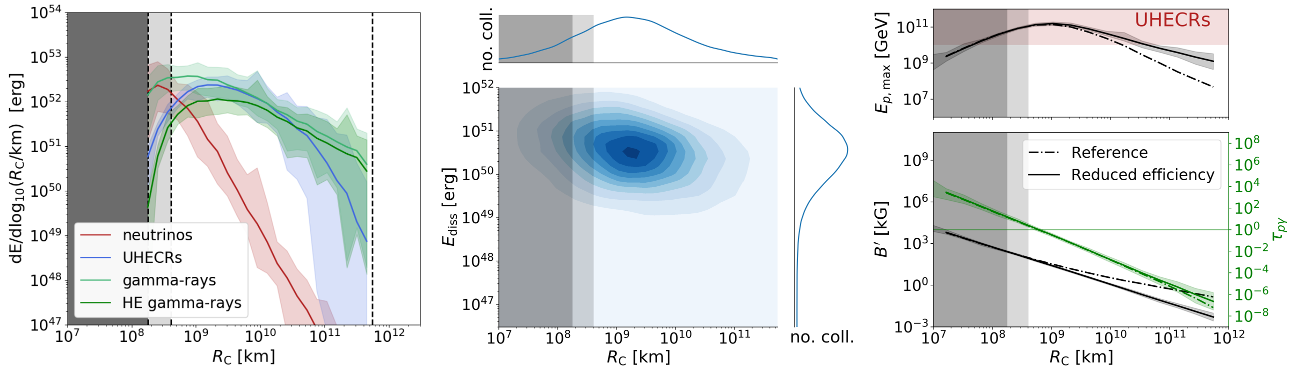

(b) Reduced efficiency model

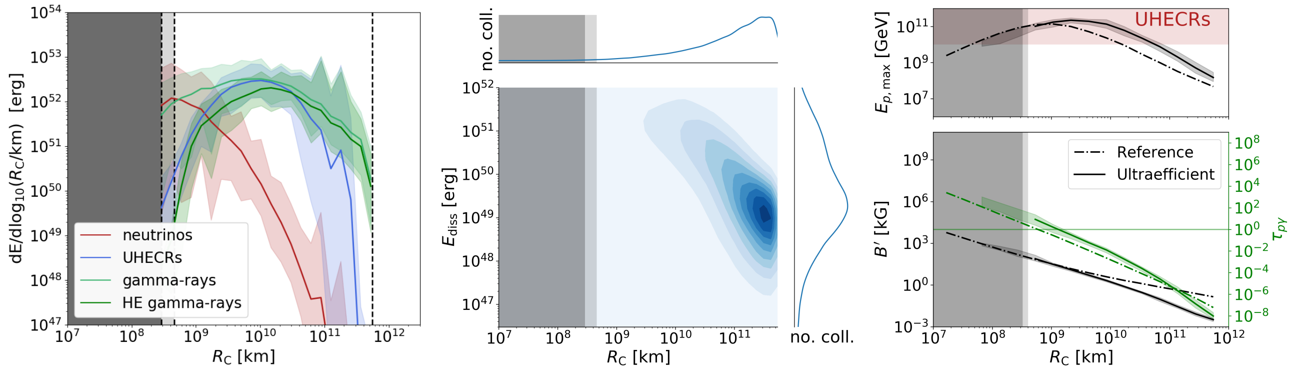

(c) Ultra-efficient model

Figure 6.11: Illustration of the fireball radiation output for different collision models. The format is the same as in Fig. 6.4 and Fig. 6.4. In contrast to Section 6.3, these were done with an initial equal mass setup. The different collision models are: (a) *Baseline with equal mass*, (b) *Reduced efficiency* and (c) *Ultra-Efficient*. The initial shell setup is purely stochastic as for GRB 1 in Fig. 6.4, but with fewer shells in the *Ultra-Efficient* model. See Section 6.5 and Tab. 6.3 for details on the models.

maximal UHECR energies are higher and more energy is released in UHECRs. The reason for this is that the merged shells expand to due to the remaining internal energy, thus shells in the later stages have larger width. The radiation density is therefore decreased for collisions in the later stage of the fireball compared to the baseline model. The initial collisions behave similar to the baseline model, such that the neutrino output is almost unchanged. Note that the overall efficiency ε decreases proportional to η by about 50%. To compensate this, the initial kinetic energy in this model is therefore slightly higher, as we normalized to the total gamma-ray output.

Fig. 6.11c shows the output of the *Ultra-Efficient* model. We choose fewer initial shells in this model, to reproduce approximately the same number of collisions. This implies a larger initial separation and width, and consequently larger collision radii. Additionally in the late stage of the fireball, the thermalized shells frequently bounce off each other producing many collisions with very low E_{diss} (middle panel). The shells at larger radii have larger widths compared to the baseline model, resulting in lower magnetic fields and larger UHECR energies. Consequently, the UHECR output is slightly increased similar to the *Reduced Efficiency* model. The neutrino output is slightly decreased, as the collisions are shifted to larger radii. The efficiency of the *Ultra-Efficient* is increased to $\varepsilon \simeq 36\%$ compared to the *Reduced efficiency* with the same dissipation per collision of $\eta = 0.5$. This efficiency is close to the baseline model with $\eta = 1.0$. Since we normalize to $E_{\gamma,\text{tot}}$, the initial kinetic energy is higher by a factor of 2 in the *Reduced efficiency* model.

Overall, the impact of the collision model is rather small. While the *Ultra-Efficient* produces many more collisions at higher radii, these contribute little to the total dissipated energy. The neutrino output is only affected within a factor of 2. While the fluctuations are higher in the *Ultra-Efficient* model, they will average out for a population of sources. Additionally, hydrodynamic simulations show that a splitting into two shells as in the *Ultra-Efficient* needs a specific setup to work, while the baseline model is more realistic for most of the parameter space. In Rudolph et al. [4], we therefore discussed a more realistic model in which each collision is computed in PLUTO resulting in a model that is qualitatively similar to the *Reduced Efficiency* model, but with higher overall efficiency closer to the baseline model. However in any case, the impact of the collision model is small compared the effect of different initial engine setups that was discussed in Section 6.4 and the choice of the initial energy distribution (equal-mass or equal-energy).

6.6 Summary and discussion

In this chapter, we explored the internal shock model of GRBs. In this multi-collision model, the outflow of the GRB jet is simulated as a sequence of discrete relativistic plasma shells with different speeds, which emit radiation in cosmic rays, gamma rays and neutrinos when they collide. In contrast to one-zone models, which consider only a single representative collision, multi-zone models consider multiple collisions with very different physical conditions.

We demonstrated with different examples how the initial setup of the central engine can produce various features of observed GRB light curves. The collisions are determined by the spatial separation and speed difference of the shells. If the initial distribution of shell speeds is 'disciplined', the initial collisions will only dissipate little energy, producing small peaks with fast time variability. Broader peaks in the light curve can instead be produced by the collisions of different regions in the fireball with different average speeds. On average, the collisions for these 'disciplined' GRBs happen at larger radii, producing relatively small neutrino fluxes. If however the initial distribution of speeds is broader, the initial collisions between neighboring shells will have high efficiency. This leads to a light curve that is dominated by fast time variability. For such GRBs, the neutrino production efficiency is always high, since the initial collisions happen at small radii with high radiations densities. The GRBs with strong variability seem more representative for the catalogue of observed long GRBs, which consists mostly of simple light curves, while the cases presented in the literature are typically those with more interesting, non-trivial features.

A general feature of the multi-collision model is that the different messengers come from different regions of the fireball: neutrinos are produced at low radii, where radiation densities are high; UHECR protons come from intermediate region, where efficient acceleration and radiation losses are balanced; and high-energy gamma rays tend to come from larger radii, where their maximal energy is not limited by pair-production. This leads to a smaller predicted neutrino flux compared to one-zone predictions. The average production region in one-zone models is more compact, since it is determined from t_{var} following Eq. (6.1).

In addition to the initial GRB engine we have tested the impact of different models of the individual collisions in the internal shock scenario. Compared to the *Baseline* model, in which all internal energy is dissipated per collision, we introduced two alternative models: A *Reduced Efficiency* model, in which a fraction of the internal energy remain in the shell and leads to adiabatic expansion and an *Ultra-Efficient* model in which the remaining internal energy leads to a splitting into two shells in the final state.

The *Ultra-Efficient* model was introduced to alleviate the efficiency problem of the internal shock model, i.e. if the internal shocks do dissipate a high enough percentage of the kinetic energy, it will lead to a high prediction for the afterglow, which is in tension with observations. While the *Ultra-Efficient* model dissipates less energy per collision, the two remaining shells lead to more subsequent collisions and therefore a higher overall efficiency ε. While this model does indeed increase the efficiency, it cannot be reconciled with hydrodynamic simulations, as the splitting in two shells requires rather specific parameters of the colliding shells [4]. Even if these are realized for the initial collisions, they will not be realistic for the full fireball evolution. Still, the collision model comparison shows that the multi messenger emission is relatively robust under details of the hydrodynamic model. The conclusions derived for different engines in the *Baseline* model for the particle emission should therefore be relatively robust.

For the different setups of the initial GRB engines we also predicted the expected neutrino fluxes for a fixed rate of 667 yr^{-1} and all equal GRBs, as it was also assumed for the IceCube GRB stacking limits [24]. For this fixed rate and a baryonic loading of 10, the expected flux is at a level of $10^{-11} - 10^{-10}\ \frac{\mathrm{GeV}}{\mathrm{cm}^2\,\mathrm{s}\,\mathrm{sr}}$, which is still below the current limit. However, the expected flux is likely within of IceCube-Gen2 [99].

Note that for this study the baryonic loading of 10 and pure proton composition were an ad-hoc choice and therefore lead to a somewhat arbitrary normalization of the neutrino flux. If GRBs are actually the sources of UHECRs the baryonic loading and composition can be deduced from a fit to the UHECR data. We only discussed the impact of such a fit qualitatively. From energy budget considerations, the UHECR output of our model with a baryonic loading of 10 will not be sufficient and the baryonic loading will have to be increased by another factor of ~ 10. Additionally, the shape of the ejected UHECR spectrum is relatively soft due to the superposition of many different collisions. The total spectrum will be even softer for the superposition of the emission from many different GRBs. An UHECR output compatible with the hard spectral index found for our generic parameter scan therefore likely requires a tuning of the assumptions for engine setup, composition and escape assumption. Still, it was demonstrated in Refs. [122, 123] for one specific setup, that the GRB multi-collision model can in principle provide a good fit to the UHECR spectrum and composition. An open question is how much fine-tuning this requires for the source parameters. By combining the technology of the UHECR fit for the generic model and the multi-collision model developed for this thesis, we therefore aim to study the impact of engine parameters on the UHECR fit in a more systematic approach. This will also give a robust prediction of the of source neutrino flux, and reveal which sensitivity is required to robustly reject GRBs as the sources of UHECRs.

The conclusion that neutrinos, UHECRs and gamma-rays come from different dissipation radii means that their fluxes are not always linearly related. Stacking analyses, such as that performed by IceCube [24], usually assume simple analytic models that scale the neutrino flux with the gamma-ray luminosity. In more realistic models the scaling with gamma-ray luminosity is more complicated. We demonstrated that the observed light curve can be used as an additional indicator of the neutrino production efficiency. This could be used in future stacking analyses to scale the neutrino emission more realistically, which might result in stronger stacking bounds as the main neutrino emission is expected from only a subset of the observed GRBs.

Chapter 7

Case study of GRB170817A: off-axis neutrino emission

In this chapter, we discuss the expected neutrino flux from GRB170817A. This specific GRB has gotten large attention in the recent literature as it was observed in coincidence with the gravitational wave event GW170817. As the first gravitational wave event with an observed electromagnetic counterpart, it constitutes a significant breakthrough in multi-messenger astronomy. GRB170817A is an outlier in the observed population of GRBs with very low luminosity, which has been explained in terms of different jet models, including an off-axis observations. We will discuss the expected neutrino fluence from this event, considering different jet scenarios.

This chapter is based on the previous publication in Biehl et al. [2]. For this study, I derived the analytic scaling relations for the off-axis neutrino flux. I also contributed to cross-checks of the numerical results and the evaluation of the parameter scan.

7.1 Observations of GW170817 and GRB170817A

On August 17th 2017 the LIGO and Virgo detectors observed the binary neutron star merger GW170817. It was the closest and most precisely localized gravitational wave event to that date. 1.7 s later, a short GRB, GRB170817A, was observed in spatial coincidence by the Fermi-GBM and INTEGRAL detectors [42]. The event was eventually also observed in the X-ray [289, 290] and radio [291, 292] bands. In total the electro-magnetic follow up campaign included over 70 different observatories. A detailed overview can be found in Abbott et al. [43]. The UV, optical, and near-infrared counterparts have been interpreted as a kilonova [293–295], which shows evidence for the synthesis of heavy elements through the r-process [296, 297]. The afterglow of the event, that has since been observed, is rather unusual. It shows an exceptionally long rise until 160 days after the event and a rapid decline thereafter [298, 299].

Fermi-GBM measured the duration of the burst $T_{90} = 2.0 \pm 0.5$ s and the minimum variability timescale $t_v = 0.125 \pm 0.064$ s [300]. With the accompanying gravitational wave signal, the redshift of the source was determined to be $z = 0.008^{+0.002}_{-0.003}$, which corresponds to a distance of roughly 40 Mpc [43]. The host galaxy of the event was identified as the nearby galaxy

NGC 4993 [301], which is consistent with the redshift. From the observed gamma-ray flux of $F_\gamma = (5.5 \pm 1.2) \times 10^{-7} \frac{\mathrm{erg}}{\mathrm{s\,cm^2}}$ measured by Fermi-GBM [300], the isotropic equivalent luminosity of the source can be estimated as $L_\gamma = 10^{46.9}$ erg s^{-1}, and a corresponding isotropic equivalent energy output of $E_\gamma = 1.6 \cdot 10^{47}$ erg. This value is exceptionally low, making GRB170187 a dim outlier in the population of observed GRBs [42].

By default Fermi-GBM uses four different spectral shapes to fit the spectral energy distribution (SED) of each observed GRB. For this event the best fit was found for a comptonized spectrum [42], which we will use in this chapter. We use the 256 ms time-integrated selection from T0-0.192s to T0+0.064s, for which the fit yields a spectral index of $\alpha = 0.14 \pm 0.59$ and a peak energy $E_{\mathrm{peak}} = 215 \pm 54$ keV [300]. The observations do not give conclusive information on the Lorentz factor Γ and the observation angle θ_{obs} of the initial prompt emission. We therefore treat them as free parameters. They can however be constrained, for example by the time delay between gravitational wave and electromagnetic signal [302, 303], as we will discuss in Section 7.2.

Follow-up searches in high-energy neutrinos by ANTARES, IceCube and Auger collaborations [304] found no coincident events in a wide energy range (100 GeV to 100 EeV). These results were discussed in terms of generic models of neutrino emission from neutron-star binary mergers [305, 306], re-scaled to the observed distance and luminosity. However, no direct prediction of the neutrino flux for this event GRB170818A had been derived previous to Biehl et al. [2].

7.2 Considered jet models and radiation model

The very low luminosity of the prompt emission of GRB170817A has been accounted to the jet geometry [42]. We will consider two of the discussed jet scenarios: A structured jet with low-luminosity and a uniform top-hat jet observed off-axis (similar to the simple and advanced models in Denton and Tamborra [307]). These two geometries are illustrated in Fig. 7.1. Cocoon emission has been discussed as another mechanism to explain the dimness of GRB170817A. However, it will not be considered here. See also Mooley et al. [308] for a more detailed discussion of possible geometries and their constraints.

When considering an off-axis observation, it is important to clearly distinguish between the Lorentz factor Γ of the jet and the Doppler factor D. The latter is used to boost from the jet rest frame to the observer's frame. The Doppler factor is given as a function of off-axis angle θ

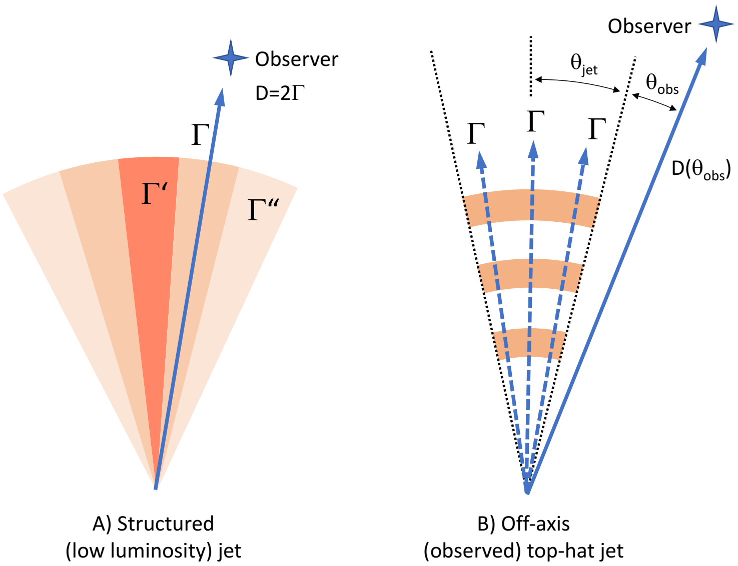

Figure 7.1: Illustration of different jet geometries and corresponding angles used in this work. Figure taken from Biehl et al. [124]

and Lorentz factor Γ by

$$D(\theta) = \frac{1}{\Gamma(1-\beta\cos\theta)} \approx \frac{2\Gamma}{1+\theta^2\Gamma^2}. \tag{7.1}$$

For a jet observed on-axis this means $D(0) = 2\Gamma$. Note that studies of a whole population of sources, such as in Chapter 6, typically assume an average observation angle of $\theta = \Gamma^{-1}$, giving $D = \Gamma$. Because of this the Doppler and Lorentz factor are not always clearly distinguished in the literature. However, in the off-axis scenario this distinction is important.

Structured jet: In this scenario the jet is assumed to be structured in the sense that it has different characteristics, such as luminosity L_γ and Lorentz factor Γ, depending on the jet axis under which it is observed (Fig. 7.1, A). The dimness of GRB170817A is then simply explained by us observing the event at an axis with lower luminosity. Note however that the observer will also be exposed to the off-axis emission from the other axes of the jet, which are illustrated by Γ' and Γ'' in the figure. While they are suppressed due to their Doppler factor, they might still overshoot the on-axis radiation if the gradient in luminosity in the jet is too high. This implies that the Γ-factor and brightness must not change too rapidly with the viewing angle. As long as these variations are small, the structured jet can be treated as a low luminosity top-hat jet observed on-axis. We will therefore also refer to this case as "low luminosity" jet.

Off-axis observation: In this scenario, we assume a uniform top-hat jet that is observed off-axis (Fig. 7.1, B). For an off-axis observer, the luminosity of the jet is suppressed due to the Doppler factor, which depends on the angle between the edge of the jet and the observation axis θ_{obs}. The details of these transformations will be discussed in Section 7.2.2. We do not impose any direct constraint on θ_{jet}, but we will see a transition in the luminosity scaling around $\theta_{\mathrm{obs}} \simeq \theta_{\mathrm{jet}}$ due to the observed geometry. While the jet-opening angle θ_{jet} is not inferable from observations, we use the estimate $\theta_{\mathrm{jet}} \simeq 1/\Gamma$ to demonstrate this transition.

In both scenarios, we assume a one-zone internal shock model in a relativistically expanding fireball. The radiation model is similar to the radiation part of the multi-collision model in Section 6.2.2, but with some collision parameters like luminosity fixed by the observations. For an on-axis observer, this fireball is not distinguishable from a sphere expanding into the full solid angle, which allows to treat the fireball in the isotropic-equivalent picture. For the low luminosity jet we can take the observed quantities at face value, while for the off-axis scenario, we have to transform the observation to the on-axis frame as a function of $(\theta_{\mathrm{obs}}, \Gamma)$.

7.2.1 One-zone model and dissipation radius

Compared to the multi-zone model in the previous chapter, we only compute a single collision for the one-zone internal-shock model. The result is then scaled to the total number of collisions estimated by the duration and variability of the burst as $N_{\mathrm{coll}} \approx T_{90}/t_{\nu}$. The critical input for the neutrino production efficiency of this collision are the gamma-ray (isotropic equivalent) luminosity L_{γ} and the collision radius R_{coll} where the shells in the jet collide, shocks form, and particles are accelerated. The collision radius and width of the radiating shell is typically estimated by the relations introduced in the previous chapter, Eq. (6.1), where $R_{\mathrm{coll}} \simeq 2\Gamma^2 c t_{\nu}$. The effect of cosmological redshift can be neglected in this case due to the small redshift under which GRB170817A was observed. Note that in this formula R_{coll} and Γ are both given in the source frame, while t_v is given by the on-axis observation. The radius R_{coll} therefore indirectly depends on the observation angle, as t_v scales with θ_{obs}

As we discussed for the multi-collision models, radiation from the emitting shell can only be observed if the shell is optically thin to Thomson scattering. The optical thickness to Thomson scattering was derived in Eq. (6.9). Using the one-zone estimates and solving for photospheric radius R_{ph} at $\tau_T = 1$, we get

$$R_{\mathrm{ph}} \simeq \left(\frac{\sigma_T}{4\pi m_p} \right)^{1/2} \left(\frac{\xi_A}{\varepsilon} \frac{E_{\mathrm{iso,on}}}{N_{\mathrm{coll}}} \right)^{1/2} \tag{7.2}$$

where $E_{\text{iso,on}}$ is the total isotropic equivalent energy in γ-rays for an on-axis observer, ε is the conversion efficiency of kinetic energy to total dissipated energy, and ξ_A is the baryonic loading defined as ratio between energy in protons and photons in the FERMI-GBM energy band from 10 to 1000 keV. The efficiency ε is not predicted by the one-zone input parameters $\langle\Gamma\rangle$ and t_ν. We therefore use use $\varepsilon = 25\%$ as an estimate corresponding roughly to the average values found in Chapter 6. The collision can only be observed if $R_{\text{coll}} \gtrsim R_{\text{ph}}$, which therefore constrains the baryonic loading. As the photospheric radius scales with the ratio ξ_A/ε, this means that lower values for ε corresponds to a lower maximal baryonic loading for fixed R_{ph}. Note that there is no lower limit on the baryonic loading from observations. While the maximal baryonic loading limits the neutrino fluence prediction, there is no guaranteed neutrino flux for this GRB.

In addition to the lower limit due to R_{ph}, the collision radius is also constrained from above by the observed time delay between the gravitational-wave and gamma-ray signals of $t_{\text{delay}} = 1.7\text{s}$ [42]. Assume that the emission originates from the collision of two shells with Lorentz factors Γ_1 and Γ_2 and $\Gamma_1 \ll \Gamma_2$. If we assume that the first shell is emitted at the same time as the gravitational wave and the second one with a time delay of t_{delay}, they will collide at a radius $R_{\text{coll}} \approx 2\Gamma_1^2 c t_{\text{delay}}$. The upper limit on the collision radius for $\Gamma_1 < 100$ is therefore around $R \sim 10^{9.5}$ km. For the measured $t_\nu = 0.125 \pm 0.064\,\text{s}$ the estimated collision radius is between 10^7 and 10^8 km, which is well below this upper limit. Additionally, the time delay imposes a limit on the observation angle. The plasma jet travels with speed $\beta \cdot c$, while the gravitational wave travels with c. If they are emitted at the same time, this alone leads to a difference in the observation of the gamma-ray and gravitational wave signal of [303]

$$t_\gamma - t_{\text{GW}} = \frac{R_{\text{coll}}}{\beta \cdot c}(1 - \beta\cos(\theta_{\text{obs}}))\,. \tag{7.3}$$

This implies only an upper limit on θ_{obs}, since the observed time delay could be partly accounted for by a delay between the neutron star merger and the jet launch.

The numerical simulation for these one-zone estimates are done using NEUCOSMA, following essentially the methods for the radiation of a single shell in the multi-collision model described in Section 6.2.2. For the target photon spectrum, we assume the observed SED as mentioned in Section 7.1. The one-zone model however needs the on-axis magnitudes such as luminosity as input. If we assume an off-axis observation, the observed values therefore have to be transformed to the on-axis frame first. The resulting fluxes then needed to be transformed back to the off-axis frame. These transformations will be discussed in the following section.

7.2.2 On-off-axis transformations

For an emitting shell moving at relativistic speed, the observed quantities such as energy and time will be Doppler shifted depending on the observation angle and Lorentz factor of the emitting shell

$$\begin{aligned} t &= D(\theta_{\mathrm{obs}})^{-1}\, t' \\ E &= D(\theta_{\mathrm{obs}})\, E' \end{aligned} \tag{7.4}$$

with the Doppler factor in Eq. (7.1). These transformations are valid for quantities such as the peak energy $E_{\gamma,\mathrm{peak}}$, which can be defined in the shell frame. They are however not necessarily valid for observed quantities that have to be integrated over the geometry of the emitting shell. Such quantities are the isotropic equivalent energy E_{iso} and variability time t_v, and the duration T_{90} which is only defined for the whole burst.

The scaling of the isotropic equivalent energy E_{iso} depends on the observed spectral flux F_ν (measured in $[\mathrm{erg\,s^{-1}\,cm^{-2}\,Hz^{-1}}]$) emitted by the radiating surface of the shell, which is integrated over observed time, area and frequency. As the observation is dominated by different regions of the surface depending on the observation angle, E_{iso} scales differently depending on whether the observer is inside or outside of θ_{jet}. A full derivation can be found in the Appendix of Ioka and Nakamura [309], arriving at:

$$E_{\mathrm{iso}}(\theta_{\mathrm{obs}}) \propto \begin{cases} \mathrm{const} & \text{for } \theta_{\mathrm{obs}} \lesssim 0 \\ D(\theta_{\mathrm{obs}})^2 & \text{for } 0 < \theta_{\mathrm{obs}} \lesssim \theta_{\mathrm{jet}} \\ D(\theta_{\mathrm{obs}} + \theta_{\mathrm{jet}})^3 & \text{for } \theta_{\mathrm{jet}} < \theta_{\mathrm{obs}} \end{cases} \tag{7.5}$$

These three regimes have different geometrical interpretations:

$\theta_{\mathrm{obs}} < 0$: The observer is in the jet-opening angle, which means that the jet looks like a spherical fireball. Most radiation comes from the angles close to the viewing axis. This is the classical on-axis case, in which the jet is not distinguishable from a sphere expanding over the full solid angle.

$0 < \theta_{\mathrm{obs}} \lesssim \theta_{\mathrm{jet}}$: The observer is outside the opening angle, but only at a small angle from the jet edge. Therefore the jet geometry still contributes, and the observed flux is to be integrated over the observable part of the jet surface close to the edge.

$\theta_{\mathrm{jet}} < \theta_{\mathrm{obs}}$: The observer sees the fireball under a larger angle. All observed regions of the fireball surface then have approximately the same Doppler factor. Therefore the jet looks like a point source to the observer.

Due to these transformations the same jet clearly looks different when observed on- or off-axis. However, in our case the observations (E_γ, E_peak, t_ν, T_{90}) are fixed. We are instead interested in what they imply for values in the shell rest frame, if they are interpreted as either on- or off-axis. To derive these relationships, we define the ratio of the on- and off-axis Doppler factors as

$$b \equiv \begin{cases} D(0)/D(\theta_\mathrm{obs}) & \text{for } 0 \leq \theta_\mathrm{obs} \lesssim \theta_\mathrm{jet} \\ a\, D(0)/D(\theta_\mathrm{obs}+\theta_\mathrm{jet}) & \text{for } \theta_\mathrm{jet} < \theta_\mathrm{obs} \end{cases} \tag{7.6}$$

with $a = D(2\theta_\mathrm{jet})/D(\theta_\mathrm{jet})$ chosen by demanding that b is continuous in θ_obs. Note that with this definition always $b \geq 1$. With Eq. (7.5), we can express the on-off-axis ratio for the isotropic equivalent energy as

$$\frac{E_\mathrm{iso,on}}{E_\mathrm{iso,off}} = \begin{cases} b^2 & \text{for } \theta_\mathrm{obs} < \theta_\mathrm{jet} \\ b^3 & \text{for } \theta_\mathrm{jet} < \theta_\mathrm{obs} \end{cases}, \tag{7.7}$$

while the peak energy is simply Doppler shifted by

$$\frac{E_\mathrm{peak,on}}{E_\mathrm{peak,off}} = b\,. \tag{7.8}$$

Similar to E_iso, the variability timescale has to be integrated over the shell geometry, as the radiation from different parts of the shell surface is delayed depending on θ_obs. This means that the off-axis observer sees a broader peak. Following Ioka and Nakamura [309] we scale it as:

$$\frac{t_{v,\mathrm{on}}}{t_{v,\mathrm{off}}} = \begin{cases} b^{-1} & \text{for } \theta_\mathrm{obs} < \theta_\mathrm{jet} \\ b^{-1/2} & \text{for } \theta_\mathrm{jet} < \theta_\mathrm{obs} \end{cases}. \tag{7.9}$$

There is a disagreement about the scaling of the duration T_{90}. Salafia et al. [310] argue that the observed burst duration T_{90} does not change depending on the observation angle because it is defined in the source frame, which is at rest relative to the observer. However, in Abbott et al. [42] it is implied that the duration of the burst T_{90} scales with b depending on the observation angle as well, as they assume a fixed number of collisions N_coll. We chose not to re-scale T_{90} with off-axis angle, which implies a larger number of collisions N_coll according to Eq. (7.9) ($N_\mathrm{coll} \simeq T_{90}/t_v$) in the on-axis frame. While this may be counter-intuitive, the physical picture is that the many peaks in the light curve observed on-axis are smeared out off-axis as discussed in Salafia et al. [310]. The observer therefore cannot distinguish all collisions, leading effectively to a smaller time variability if observed off-axis. The quantitative effect is small anyway, as the number N_{coll} drops out from the computation to a first approximation, since the injected L_γ is divided by N_{coll} while the neutrino flux is scaled up by N_{coll}. The smaller value of t_v in the on-axis frame only slightly increases the neutrino production efficiency because it is used to estimate the shell width and collision radius. We have also checked numerically that the impact of re-scaling T_{90} does not change the qualitative picture.

7.2.3 Scaling of the neutrino flux with observation angle

The relationships reported in Section 7.2.2 describe the relation between on- and off-axis observables, leading to e.g. lower observed energies for an off-axis observer if the on-axis observables are fixed. Conversely, if the off-axis observables are fixed by observations, higher energies and shorter timescales are obtained in the shell frame. Note that the on-axis observables still have to be boosted to the shell frame with the Doppler factor $D(0)^{-1} = 1/(2\Gamma)$. While the secondary radiation in the shell rest frame is higher for the off-axis interpretation, it is then reduced by the boost back to the off-axis observer. Still, there is a net effect leading to a higher neutrino flux for larger θ_{obs}. In the following analytical discussion, we focus on the case of small angles $\theta_{\mathrm{obs}} < \theta_{\mathrm{jet}}$ for the off-axis transformations. Quantities in the shell frame are expressed by primed quantities (e.g. E').

The gamma-ray peak in the shell frame is shifted to higher energies as $E'_{\gamma,\mathrm{peak}} \propto b$, which implies that the neutrino production threshold is lower. As $E'_{\nu,\mathrm{peak}} \propto 1/E'_{\gamma,\mathrm{peak}}$ (higher gamma energies lead to lower production thresholds), the peak of the observed neutrino spectrum will scale with $E_{\nu,\mathrm{peak}} \propto b^{-2}$. The neutrino production efficiency f_ν (the energy fraction the protons dumped into neutrino production) scales with the particle densities in the shell which depend on the luminosity and the dissipation radius. It can be estimated for the on-axis case from the pion production efficiency as [121, 275]

$$f_\nu \equiv \frac{E_{\nu,\mathrm{iso}}}{\xi_A E_{\gamma,\mathrm{iso}}} \propto \frac{L_\gamma}{\Gamma^4 E_{\gamma,\mathrm{peak}} t_v}\,, \tag{7.10}$$

if the synchrotron cooling of the secondaries is neglected. For small angles $\theta_{\mathrm{obs}} \leq \theta_{\mathrm{jet}}$, the product $E_{\mathrm{peak}} t_v$ is invariant under the observation angle and f_ν transforms proportional to luminosity $L_\gamma = E_{\gamma,\mathrm{iso}}/t_\nu$ as $f_\nu \propto b^3/\Gamma^4$.

As the ratio $E_{\nu,\mathrm{iso}}/E_{\gamma,\mathrm{iso}}$ is invariant under observation angle while the neutrino peak is shifted by $\propto b^{-2}$, an approximate scaling of the observed neutrino fluence F_ν in [$\mathrm{GeV}^{-1}\mathrm{cm}^{-2}$] can be obtained as

$$F_{\nu,\mathrm{off}}(E_\nu) \approx \frac{b^5}{\Gamma^4} F_{\nu,\mathrm{on}}(b^2 E_\nu)\,. \tag{7.11}$$

This means that the expected neutrino flux is enhanced for fixed baryonic loading if the observation is interpreted as off-axis- rather than on-axis-emission, while it is also shifted to lower energies.

Since the baryonic loading is limited by the photospheric radius in Eq. (7.2), the scaling of $E_{\mathrm{iso}} \propto b^2$ has an additional implication for baryonically loaded jets. As the baryon and electron densities scale with the energy density, the shells are more opaque to gamma rays if observed off-axis with the same gamma-ray flux. The scaling can be read off from Eq. (7.2), where only $E_{\gamma,\mathrm{iso,on}} \propto b^2$ is to be re-scaled:

$$R_{\mathrm{ph}} \simeq \left(\frac{\sigma_T}{4\pi m_p}\right)^{1/2} \left(\frac{\xi_A}{\varepsilon} \frac{E_{\gamma,\mathrm{iso,off}}}{N_{\mathrm{coll}}}\right)^{1/2} \left(\frac{b^2}{\Gamma}\right)^{1/2} \tag{7.12}$$

This condition effectively limits θ_{obs} and Γ for a fixed baryonic loading if the emission ought to come from the dissipation in internal shocks beyond the photosphere. It can be used to estimate the maximal allowed baryonic loading for which the shell is still transparent at R_{coll} from the condition $R_{\mathrm{coll}} \gtrsim R_{\mathrm{ph}}$ as

$$\frac{\xi_A}{\varepsilon} \lesssim \frac{4\pi m_p}{\sigma_T} \frac{N_{\mathrm{coll}}}{E_{\mathrm{iso,\ off}}} \frac{4\Gamma^5}{b^4} \tag{7.13}$$

For larger baryonic loadings, the radius where internal shocks develop will be below the photosphere – where gamma-rays cannot escape. This is therefore the maximal baryonic loading allowed for the internal shock model. The neutrino flux computed for this value then also corresponds to the maximal allowed flux within this model.

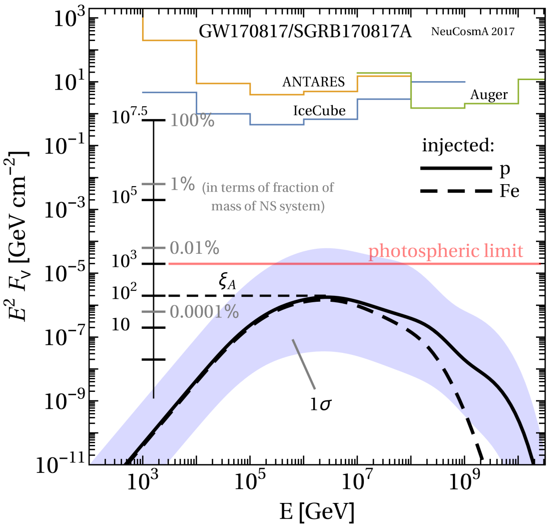

Figure 7.2: Fluence of $\nu_\mu + \bar{\nu}_\mu$ for GRB170817A assuming pure proton (solid) or iron (dashed) injection. The input parameters are $z = 0.008$, $L_X \sim 10^{47}\,\mathrm{erg\,s^{-1}}$, $t_v = 0.125\,\mathrm{s}$, as given in Abbott et al. [43], Goldstein et al. [300], and $\Gamma = 30$. The 1σ-region includes the uncertainty of the measurement of these parameters given in Section 7.1. The black scale indicates how the fluence will change with the baryonic loading ξ_A, with the grey percentage representing the fraction of the total mass of the NS merger. For a baryonic loading greater than $\xi_A = 10^3$, this collision would be sub-photospheric according to Eq. (7.13), indicated by the horizontal red line. Neutrino limits are taken from Albert et al. [304]. Figure taken from Biehl et al. [2].

7.3 Neutrino flux predictions

First, we assume a structured (low luminosity) jet observed on-axis as described in Section 7.2, where the Lorentz factor is fixed to $\Gamma = 30$. Fig. 7.2 shows the predicted fluence of muon neutrinos for pure proton injection modeling GRB170817A. As shown already for the multi-zone model in Section 6.4.1, injecting nuclei heavier than protons in GRBs shifts the cutoff of the neutrino fluence to lower energies, while there is only a slight impact at the peak. The effect is actually smaller here compared to the multi-zone model, since we only compute a single collision at an intermediate radius, for which both elements reach the same maximal rigidity. The solid curve has been computed with an initial baryonic loading of $\xi_A = 100$. As indicated by the scale on the left side of the plot, the flux level scales directly with ξ_A, as we assume that the SED is fixed to the observations. If the SED was computed self-consistently, its shape would likely be affected by high values of ξ_A, which would also change the neutrino flux prediction. The blue band includes the 1σ-uncertainties on the measured duration T_{90}, time variability t_v, redshift z,

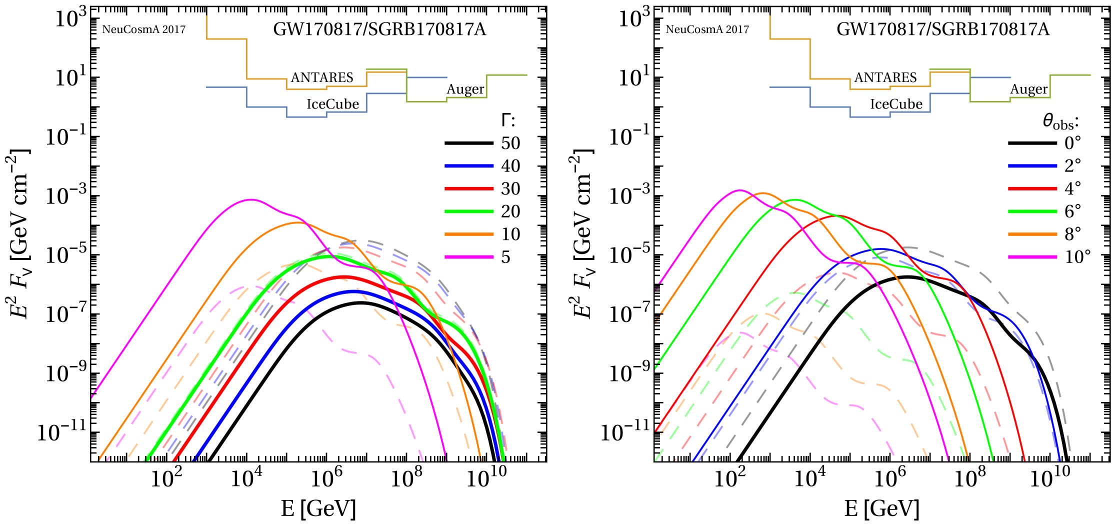

(a) Dependence on Γ (b) Dependence on $\theta_{\rm obs}$

Figure 7.3: Fluence of $\nu_\mu + \bar{\nu}_\mu$ for GRB170817A and with: (a) different values of the Lorentz factor Γ in the structured jet case with fixed $\theta_{\rm obs} = 0$; (b) different observations angles $\theta_{\rm obs} = 0$ in the off-axis scenario for fixed $\Gamma = 30$. We assume pure proton injection and the same parameters as given in Fig. 7.2. Solid curves refer to a fixed baryonic loading of $\xi_A = 100$, where thick solid curves correspond to collisions above the photosphere, and thin curves indicate sub-photospheric collisions. For the dashed curves, the baryonic loading has been maximized demanding that $R_{\rm coll} > R_{\rm ph}$. Figures taken from Biehl et al. [2].

gamma-ray fluence F_γ as well as the spectral index α and peak energy $E_{\rm peak}$ of the SED. The grey scale on the left shows fraction of the total mass $(2.74^{+0.04}_{-0.01} M_\odot$ [46]) of the binary system translated into accelerated baryons. The upper limit is given by the photospheric constraint at a baryonic loading of $\xi_A \simeq 10^3$. Even when upscaled to this extreme value, the neutrino flux prediction is still about four orders of magnitude below the sensitivity of the neutrino detectors.

We now demonstrate the scaling of the neutrino flux for the two jet scenarios.

Strutured jet: We show the impact of the Lorentz factor on the muon neutrino fluence in Fig. 7.3a. The illustrate the qualitative scaling with Γ, we show the solid curves for fixed baryonic loading $\xi_A = 100$ ignoring the photospheric constraint. The scaling agrees well with the analytic prediction from Eq. (7.11). For small Γ there is an additional damping at the high energy tail of the spectrum, which is due to secondary cooling that was ignored in the analytical estimate. As the collision radius decreases with Γ, the neutrino production becomes more efficient. However, smaller Γ also implies larger radiation densities in the shell frame, which increases the photospheric radius. The dashed curves show the maximal neutrino flux that is compatible with the photospheric constraint in Eq. (7.13). For $\Gamma < 20$ this constraint means the initial prediction with $\xi = 100$ has to be down-scaled. For lower values of Γ the neutrino flux is therefore actually more constrained than for higher values. In any case, the predictions are at least four orders of magnitude below the neutrino-telescope sensitivities in the structured jet scenario.

Off-axis fireball: In the off-axis fireball scenario, the observation angle θ_{obs} enters as an additional parameter influencing neutrino production and photospheric radius. The dependence of the neutrino fluence on the observation angle is shown in Fig. 7.3b with the Lorentz factor fixed to $\Gamma = 30$. The solid curves are again given for fixed $\xi_A = 100$ and scale in agreement with Eq. (7.11). For this fixed ξ_A the collisions already become sub-photospheric for small observation angles $\theta_{\text{obs}} \sim 2°$. For large observation angles, the fluence is therefore highly suppressed by the photospheric constraint. This is again indicated by the dashed lines, which are scaled to the maximal ξ. Compared to the structured low luminosity jet, the off-axis scenario therefore makes it even less likely to detect a neutrino from this event.

As demonstrated for the two scenarios, small Lorentz factor Γ or large observation angle θ_{obs} imply a strong photospheric constraint on the baryonic loading of the GRB jet. In Fig. 7.4a we show how this constraint affects the full parameter space. The red contours show the maximal baryonic loading allowed such that the collision is still super-photospheric in the internal shock model. The scaling with θ_{obs} in Eq. (7.6) has a break at $\theta_{\text{obs}} = \theta_{\text{jet}}$, which is indicated by the white dashed line assuming that $\theta_{\text{jet}} = 1/\Gamma$.

The observed time delay of $t_{\text{delay}} = 1.7$ s between gravitational wave and gamma-ray signal implies a constraint on the observation angle [303], as explained in Section 7.2.1. The allowed region is highlighted in white.[1] Note that this is only an upper limit, which assumes that the gravitational wave and jet are emitted at the same time. If there is already a time delay between the inital collision and the jet launch, the time delay due to the observation angle can be smaller. However, the stronger constraint comes from the photosphere. A baryonic loading $\xi_A = 10$ is a common ad-hoc choice in the literature for the neutrino production from GRBs. For this value the photospheric limit is already more constraining than the time delay.

The uncertainty on the neutrino flux implied by the allowed parameter space is shown in Fig. 7.4b. The blue band shows the range of the neutrino flux allowed within the parameter

---

[1] Compared to the calculation given in Salafia et al. [303], we assume a different efficiency $\varepsilon = 0.25$ to convert between the kinetic energy and the isotropic equivalent energy, leading to a slightly larger allowed region.

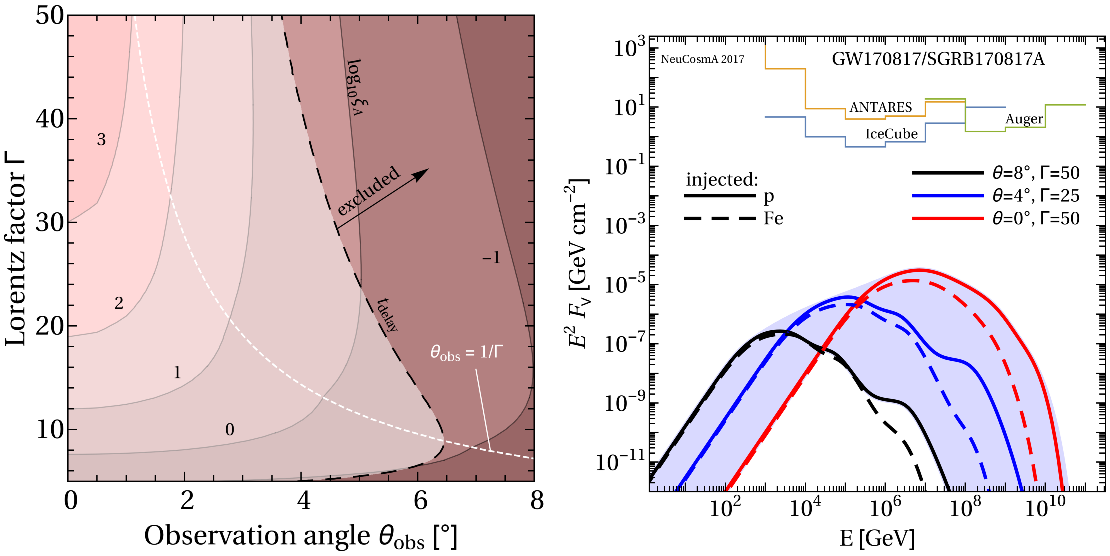

(a) Parameter space

(b) Neutrino flux range

Figure 7.4: Allowed parameter space in $\theta_{\mathrm{obs}} - \Gamma$ for GRB170817A and resulting neutrino fluence: (a) The red contours in $\theta_{\mathrm{obs}} - \Gamma$ show the maximal baryonic loading $\xi_{A,\mathrm{max}}$ such that the dissipation radius is still super-photospheric. The scaling in Eq. (7.7) and (7.9) has a break at $\theta_{\mathrm{obs}} = \theta_{\mathrm{jet}}$, which is illustrated by the white dashed curve assuming that $\theta_{\mathrm{jet}} = 1/\Gamma$. The dashed black curve corresponds to a 1.7 s photon arrival time delay between gamma-ray and gravitational-wave signal according to Salafia et al. [303]. The dark shaded region would lead too large arrival-time delays, and is therefore excluded. (b) The blue band shows the range of the $\nu_\mu + \bar{\nu}_\mu$ fluence for the allowed parameter space in $\theta_{\mathrm{obs}} - \Gamma$ assuming pure proton injection and the maximal baryonic loading. Three example curves are shown with fixed parameters (see legend) and pure proton (solid) or pure iron (dashed) injection. Figures taken from Biehl et al. [2].

space $\Gamma - \theta_{\mathrm{obs}}$, assuming that for each combination the maximal baryonic loading is realized. For illustration we show three discrete parameter combinations for proton (solid) and iron (dashed) injection. As indicated already by Fig. 7.3, the the neutrino fluence will not exceed $\sim 5 \times 10^{-5}$ GeV cm^{-2}, which is about a factor 10^{-4} below the sensitivity of the neutrino telescopes. It is therefore highly unlikely to see any neutrinos produced in the prompt phase of this short GRB, under the assumption that the prompt gamma-ray and neutrino emission come from the same dissipation radius in the internal shock model.

Note again that rescaling with the maximum baryonic loading means that we do not allow for sub-photospheric collisions. If the photospheric constraint was omitted, i.e. by a model considering sub-photospheric emission, the neutrino peaks could be strongly increased. The reason is that for smaller radii the energy density would be much higher resulting in highly efficient neutrino production. However, Fig. 7.3 indicates that even for the extrapolation of our model below the photosphere the peak is well below the sensitivity of the neutrino telescopes. On the other hand the baryonic loading is not constrained from below, as the jet might be purely leptonic. There is therefore no lower limit on the neutrino flux.

7.4 Discussion

We have computed the expected neutrino fluence from GRB170817A for two different scenarios: a structured jet scenario, and an off-axis top-hat scenario. The flux is strongly constrained by the baryonic loading, since larger values of the baryonic loading imply strong photospheric constraints. Still, even without these constraints the flux is far below the sensitivities of the neutrino searches. This is not unexpected, since there was no neutrino candidate associated with GRB170817A. Nevertheless, we demonstrated that the scaling with the off-axis angle is more complicated than the scaling assumed in the initial neutrino searches [304]. This should be taken into account for possible future off-axis observations.

Note that we only discussed two rather simple jet scenarios and used only approximate scaling relations for the off-axis observation, that transfer them to on-axis quantities. In a recent study, Ahlers and Halser [311] derived the scaling for a more detailed treatment of the off-axis observation. They calculate the full integral of the neutrino flux over the observed emitting shell for both a top-hat and structured jet. This more detailed method leads to slightly higher neutrino fluxes in the off-axis case, see their Fig. 4. The flux prediction is, however, still far out of reach of the detection limits for this event.

In fact, the jet scenario is still under discussion. Since the prompt emission from GRB170817A was rather typical for a GRB except for its dimness, the structured or off-axis jet scenarios were initially favoured over other mechanisms, such as cocoon emission [308]. However, the afterglow that has since been observed is rather untypical with an exceptionally long rise and a rapid decay [298, 299]. From these observations, Mooley et al. [308] conclude that the off-axis scenario is unlikely, since it would produce a sharper rise. Instead, they find that cocoon emission from a chocked jet is most consistent with the gamma-ray, x-ray and radio emission. However, other models find a consistent description of the afterglow for an off-axis jet scenario [312–315] or for the combination of an off-axis jet with cocoon emission [316].

Chapter 8

Conclusion and Outlook

In this thesis, we have treated different aspects of modelling the UHECR spectrum and composition: We implemented a new code, PriNCe, for the extragalactic propagation of UHECRs; we employed this code in an extensive fit to the observed UHECR spectrum and composition treating several model uncertainties, such as hadronic-interaction models and disintegration models; we also discussed a multi-collision model of GRBs and its neutrino and UHECR output. In all of theses cases, we considered the UHECR-neutrino connection, which means that high-energy astrophysical neutrinos are likely produced by photo-hadronic interactions of UHECRs and therefore carry information on their acceleration sites and their redshift evolution. Detailed summaries were given at the end of each chapter, here we recapitulate only the main results:

PriNCe: We explained the main methods employed for the propagation code PriNCe, that was developed from the ground up in the context of this thesis. PriNCe solves the integro-differential equation system for extragalactic UHECR propagation using a vectorized formulation. This leads to high computational efficiency with computation times of around half a minute for a spectrum containing iron injection. Compared to Monte-Carlo based methods, this approach is especially efficient for variations of the extragalactic environment and interaction cross sections. The code was written in a modular way, such that it can be extended in the future to include additional physical processes, for example the electromagnetic cascade and diffusion in magnetic fields. It should also be applicable to isotropic source models with relatively little extensions.

UHECR parameter scan: We presented an extensive parameter scan of a fit to the Auger energy spectrum and composition, which is the first practical application of PriNCe. We discussed the correlations between the main model parameters for a generic source with maximal rigidity R_{max}, spectral index γ and source evolution parameter m. The main features are an overall low maximal rigidity and an anti-correlation between γ and m. Both of these features are affected by the X_{max} based composition measurements, which allow for relatively low superposition of individual mass spectra. We put specific focus on the impact of different photo-disintegration and air-shower models. Especially the choice of the hadronic interaction model has a strong qualitative impact on the allowed parameter space in the fit. From the allowed UHECR parameter space, we computed the expected flux of cosmogenic neutrinos. This flux is very low independent of the exact model choices and will likely be out of reach of even future radio-based neutrino detectors.

GRB multi-collision models: We discussed GRBs in the internal shock model with multiple collisions. In this model, the GRB jet is represented by a sequence of discrete plasma shells that collide at different radii depending on their initial separation and relative speeds. In this sophisticated model, the different messengers come from different radii: Neutrinos are produced close to the photosphere, where their production efficiency is high; UHECRs escape at intermediate radii, where acceleration is efficient and not limited by interaction losses; high-energy gamma rays come from large radii, where their maximal energy is less constrained by interactions. Within this model, we discussed the impact of two assumptions: The initial engine setup and the model for the collision of two shells.

The initial engine setup can be roughly categorized into two classes: A ‘disciplined’ engine for which the initial spread in Lorentz factors is low and a ‘stochastic’ engine with large initial fluctuations in Lorentz factors. In the ‘disciplined’ case, the average collision radii are generally larger. For stochastic engines, the neutrino flux level is relatively robust under the exact setup of shells, while the disciplined case can produce significantly lower neutrino fluxes. The observed light curve is indicative of the initial shell distribution and might therefore be used as an indicator for the expected neutrino production. We also explored different models for the collision of two shells, which produce either one or two shells in the final state. The model with two final shells leads to a higher overall conversion efficiency from kinetic to dissipated energy. However, such a two-shell final state requires fine-tuned parameters of the initial shells. Assuming a single merged shell in the final state is therefore the more realistic assumption. In any case, the neutrino and UHECR output is relatively robust under a variation of the collision model.

We did not fit the result of this multi-collision GRB model to UHECR spectrum and composition. However, a qualitative discussion of the ejected UHECR spectra in our examples reveals that the spectral index is too soft to fit the observations after propagation. By adjusting the engine parameters, a good fit to the UHECR spectrum can in principle be found as previously demonstrated in Globus et al. [122]. Still, it is an open question how much fine-tuning this requires.

Neutrino flux from GRB170817A: GRB170817A has received large attention in the recent literature, as it was the first GRB observed in coincidence with a gravitational wave signal. It is an outlier in the population of observed GRBs due to its dimness, which might be explained by an off-axis observation. We computed the neutrino flux from this event under the assumption of different jet geometries and discussed their photospheric constraints. For all cases, the expected neutrino flux is at least four orders of magnitude below the experimental sensitivities, which is consistent with the non-observation of neutrino candidates from its direction. Still, our study shows some of the subtleties that need to be taken into account in the case of future off-axis observations of GRBs.

The three parts of UHECR models (source, propagation and air shower) are strongly connected. This is demonstrated by comparing the result of our UHECR fit with a generic source model to the output of the GRB multi-collision model. The generic source model leads to a good fit of the observed spectrum. Still, the resulting hard spectral indices are not easily explained by the more realistic GRB model, since its spectra will generally be softened by the superposition of individual collisions. This effect will only be amplified for a population of sources with varying parameters. The required composition ejected from the source is constrained by the observations, but varies with disintegration model and source evolution. This adds constraints on the source parameters, which should produce such a composition. A complete model of the ejected composition needs to take into account the initial abundance of elements, the acceleration mechanism, and disintegration inside the source. Additionally, the choice of air shower model changes the inferred composition significantly, which has non-trivial impact on the parameters of the generic source model and by extension also the parameters of a more realistic source model.

Often theoretical models focus on one of the three parts, such as the modelling of sources, while making ad-hoc choices about the others parts. This is required as a first assumption to focus on a detailed understanding of each part of UHECR physics. However, it is important to be aware of the effects implied by the individual model choices. The detailed connections are best studied in combined source-propagation models. With the codes and models developed in this thesis, all the tools required are present for a combined fit of the multi-collision GRB model, taking into account a wide range of model uncertainties. Judging from the qualitative discussion of the UHECR spectra, the multi-collision model parameter space will likely be strongly constrained by the fit to UHECR data. This fit will also result in a more robust normalization of the source neutrino flux than the ad-hoc choices used in Chapter 6. With this, we will be able to make a robust prediction of the neutrino flux, which is likely to be testable with next generation neutrino telescopes. This project has already started. However, conclusive results were not obtained before the completion of this thesis.

Source and propagation models can be further constrained taking into account additional data. Cosmogenic photons can limit the UHECR parameter space, but a realistic prediction has to account for the electromagnetic cascade. For models of the UHECR spectrum across the ankle, the diffusion in extragalactic magnetic fields becomes more impactful. Including these effects in PriNCe is therefore important to allow for more complete models in the future. Especially the application to source models will open up a wide range of applications.

Even without these extensions, there are a few applications of PriNCe that have not been explored yet. The code is especially efficient when varying cross sections and background photon fields. These effects have been discussed qualitatively in Alves Batista et al. [210], but only for discrete example cases. Using PriNCe, we could make more comprehensive studies giving robust uncertainty bands. Such a study would be especially interesting for photo-disintegration models, for which measurements exit only for a few isotopes and channels [18]. Additional measurements are feasible with current accelerators. Robustly identifying the cross section channels that have the strongest impact on UHECR observations will be of interest for possible future experimental efforts. Another application for PriNCe would be backward propagation of UHECRs. It is impossible to make inverted time steps in Monte-Carlo codes, since they trace the cascade from a single initial particle. However, our approach solves the differential equation for the average particle densities and can in principle be run for inverted time steps. It can therefore be used to propagate the observed flux and its uncertainties back to the assumed sources to get the allowed range for source spectra. Such an approach has been attempted in Ptuskin et al. [317] for proton and iron injection, but without simulating the full nuclear cascade, resulting in rather unphysical source spectra. Still, this approach could be explored further using our more sophisticated code. This approach needs to take into account additional assumptions, since for example the source distribution is still a free parameter. To get more realistic results, one therefore possibly needs to take into account more physical constraints, such as from secondary neutrino and gamma-ray production.

Even for more complete and realistic models of UHECR sources, a strong degeneracy between different source classes is likely to remain when considering only the diffuse UHECR data. The presented results for the UHECR fit demonstrate strong degeneracies between different source evolutions. There is also a wide range of possible UHECR source candidates with different redshift evolutions that are in principle luminous or numerous enough to power the flux of UHECRs. Source models that can in principle fit the flux and composition of UHECR include AGNs [119, 120], high luminosity GRBs [122–124], low luminosity GRBs [133, 134], TDEs [129–131], starburst galaxies [127] and galaxy clusters [163]. The discrimination between the detailed source physics from UHECR data is especially difficult since the data seems to favor a rigidity dependent maximal energy, i.e. low disintegration inside the sources, which minimizes the impact of the detailed radiation physics inside the source. One way to break this degeneracy is to include constraints from other messengers. For example, source neutrinos are strongly constraining high-luminosity GRB models [124]. However, as we demonstrated, the expected flux of cosmogenic neutrinos is very low. It will therefore likely not be constrained by observations in the foreseeable future. Another approach is to include directional information. This information is analysed in anisotropy studies for UHECRs [14]. However, these studies often consider simple, linear relations between the emission of photons and UHECRs and do not account for the energy losses. More sophisticated future studies should consider information on energy, composition and arrival direction at the same time. It has been shown that UHECR anisotropy studies are already more constraining when including energy losses for pure proton injection in a simple analytic model [318]. Such studies often rely on Markov-chain Monte-Carlo methods, which are not trivial to parallelize. They therefore require fast computation times, and hence often rely on analytic models. This might be another application for an efficient code like PriNCe, which could in principle be run inside a Markov chain.

Multi-messenger astronomy is a young field, which has only started to show its possibilities. The multi-messenger observations of GRB170817A and TXS0506+056 have been the start of what might become the 'era of multi-messenger astronomy'. A large experimental effort is going into coordinating different experiments through networks like AMON [58, 59], while also the sensitivities and spatial reconstruction will improve with future experiments like CTA [319], IceCube-Gen2 [99], and the future runs of Advanced LIGO [40]. We are therefore likely to see multi-messenger observations much more frequently in the future. This will not only allow to study the physics of the observed source in greater details, but also improve the search for correlations with populations of UHECR source candidates. A multi-messenger observation directly involving UHECRs is however not expected, since UHECRs lose spatial and temporal correlation in magnetic fields. Still, the information on the UHECR spectrum, composition and anisotropy will be improved by the future upgrades AugerPrime [85] and TA×4 [87]. Future radio detectors like GRAND [103] and ARIANNA [102] might detect neutrinos at energies above the current sensitivity of IceCube. At the same time, the microphysics of particle interactions in both photo-disintegration and hadronic interactions still constitutes a large model uncertainty. More detailed measurements of nuclear disintegration cross sections are experimentally accessible. The nuclear physics community could therefore greatly help in the modelling of UHECRs. The same is relevant for hadronic interactions. While the energies of the initial UHECR interactions in the atmosphere are not accessible in current accelerators, more detailed measurements at lower energies will help to improve model extrapolations. The observed muon excess is a clear sign of a systematic error in current hadronic interaction models [218, 219]. AugerPrime will enable a more clear separation of the electron and muon shower components, providing another tool for tuning the models. Advancements of hadronic interaction models will have crucial impact on the interpretation of UHECR composition. Especially if the shower-to-shower fluctuations

$\sigma(X_{\mathrm{max}})$ decrease in future models, they might have strong impact on the preferred UHECR source candidates as this would allow to deviate from a strictly rigidity-dependent cutoff.

The expected abundance of data from future experiments provides both a huge challenge and opportunity for multi-messenger theory. Drawing connections between the different observations requires sophisticated models that take into account detailed descriptions of the underlying microphysics and at the same time allow for efficient exploration of large parameter spaces. This includes modelling a full population of sources instead of treating a single model as representative for a whole population. To constrain the complexity of such models, they have to take into account as much of the available data as possible. The field will therefore greatly benefit from open data policies, which some experiments, such as IceCube, already employ. At the same time, making computational models open source can also facilitate the exchange. CRPROPA is an example for an open-source software that has seen large popularity and a rapid development in the recent years. While we did not have the resources to get the PRINCE code ready for publication during this PhD project, we plan to make it openly available in the near future.[1] Open-source software has the potential to greatly accelerate the progress on theoretical models and also allows various communities to use more sophisticated models in their analyses. With such combined efforts between experiments and theory, we will surely see further progress in the field of multi-messenger astronomy. This has the potential to finally reveal the origin of UHECRs and astrophysical neutrinos.

---

[1] Between the submission and publication of this thesis, the code has been published.
The source code is available at `https://github.com/joheinze/PriNCe`

Bibliography

[1] Mauricio Bustamante, Kohta Murase, Walter Winter, and Jonas Heinze. Multi-messenger light curves from gamma-ray bursts in the internal shock model. *Astrophys. J.*, 837(1):33, 2017. doi: 10.3847/1538-4357/837/1/33.

[2] Daniel Biehl, Jonas Heinze, and Walter Winter. Expected neutrino fluence from short Gamma-Ray Burst 170817A and off-axis angle constraints. *Mon. Not. Roy. Astron. Soc.*, 476(1):1191–1197, 2018. doi: 10.1093/mnras/sty285.

[3] Jonas Heinze, Anatoli Fedynitch, Denise Boncioli, and Walter Winter. A new view on Auger data and cosmogenic neutrinos in light of different nuclear disintegration and air-shower models. *Astrophys. J.*, 873(1):88, 2019. doi: 10.3847/1538-4357/ab05ce.

[4] Annika Rudolph, Jonas Heinze, Anatoli Fedynitch, and Walter Winter. Impact of the collision model on the multi-messenger emission from Gamma-Ray Burst internal shocks. *Astrophys. J.*, 893(1):72, 2020. doi: 10.3847/1538-4357/ab7ea7.

[5] Jonas Heinze, Denise Boncioli, Mauricio Bustamante, and Walter Winter. Cosmogenic Neutrinos Challenge the Cosmic Ray Proton Dip Model. *Astrophys. J.*, 825(2):122, 2016. doi: 10.3847/0004-637X/825/2/122.

[6] Alexander Aab et al. The Pierre Auger Cosmic Ray Observatory. *Nucl. Instrum. Meth.*, A798:172–213, 2015. doi: 10.1016/j.nima.2015.06.058.

[7] Darko Veberic, editor. *The Pierre Auger Observatory: Contributions to the 35th International Cosmic Ray Conference (ICRC 2017)*, 2017. arXiv:1708.06592.

[8] Michael Unger. Highlights from the Pierre Auger Observatory. *PoS*, ICRC2017:1102, 2018. doi: 10.22323/1.301.1102. [35,1102(2017)].

[9] T. Abu-Zayyad et al. The surface detector array of the Telescope Array experiment. *Nucl. Instrum. Meth.*, A689:87–97, 2013. doi: 10.1016/j.nima.2012.05.079.

[10] H. Tokuno et al. New air fluorescence detectors employed in the Telescope Array experiment. *Nucl. Instrum. Meth.*, A676:54–65, 2012. doi: 10.1016/j.nima.2012.02.044.

[11] John Matthews. Highlights from the Telescope Array Experiment. *PoS*, ICRC2017:1096, 2018. doi: 10.22323/1.301.1096.

[12] Alexander Aab et al. Observation of a Large-scale Anisotropy in the Arrival Directions of Cosmic Rays above 8×10^{18} eV. *Science*, 357(6537):1266–1270, 2017. doi: 10.1126/science.aan4338.

[13] R. U. Abbasi et al. Search for EeV Protons of Galactic Origin. *Astropart. Phys.*, 86:21–26, 2017. doi: 10.1016/j.astropartphys.2016.11.001.

[14] Alexander Aab et al. An Indication of anisotropy in arrival directions of ultra-high-energy cosmic rays through comparison to the flux pattern of extragalactic gamma-ray sources. *Astrophys. J.*, 853(2):L29, 2018. doi: 10.3847/2041-8213/aaa66d.

[15] Luis A. Anchordoqui. Ultra-High-Energy Cosmic Rays. *Phys. Rep.*, 801:1–93, 2019. doi: 10.1016/j.physrep.2019.01.002.

[16] Rafael Alves Batista et al. Open Questions in Cosmic-Ray Research at Ultrahigh Energies. *Front. Astron. Space Sci.*, 6:23, 2019. doi: 10.3389/fspas.2019.00023.

[17] A. Aab et al. Depth of maximum of air-shower profiles at the Pierre Auger Observatory. II. Composition implications. *Phys. Rev.*, D90(12):122006, 2014. doi: 10.1103/PhysRevD.90.122006.

[18] Denise Boncioli, Anatoli Fedynitch, and Walter Winter. Nuclear Physics Meets the Sources of the Ultra-High Energy Cosmic Rays. *Sci. Rep.*, 7(1):4882, 2017. doi: 10.1038/s41598-017-05120-7.

[19] Pedro Abreu et al. Interpretation of the Depths of Maximum of Extensive Air Showers Measured by the Pierre Auger Observatory. *JCAP*, 1302:026, 2013. doi: 10.1088/1475-7516/2013/02/026.

[20] Karl-Heinz Kampert and Michael Unger. Measurements of the Cosmic Ray Composition with Air Shower Experiments. *Astropart. Phys.*, 35:660–678, 2012. doi: 10.1016/j.astropartphys.2012.02.004.

[21] M. G. Aartsen et al. Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector. *Science*, 342:1242856, 2013. doi: 10.1126/science.1242856.

[22] M. G. Aartsen et al. First observation of PeV-energy neutrinos with IceCube. *Phys. Rev. Lett.*, 111:021103, 2013. doi: 10.1103/PhysRevLett.111.021103.

[23] M. G. Aartsen et al. The contribution of Fermi-2LAC blazars to the diffuse TeV-PeV neutrino flux. *Astrophys. J.*, 835(1):45, 2017. doi: 10.3847/1538-4357/835/1/45.

[24] M. G. Aartsen et al. An All-Sky Search for Three Flavors of Neutrinos from Gamma-Ray Bursts with the IceCube Neutrino Observatory. *Astrophys. J.*, 824(2):115, 2016. doi: 10.3847/0004-637X/824/2/115.

[25] M. G. Aartsen et al. Extending the search for muon neutrinos coincident with gamma-ray bursts in IceCube data. *Astrophys. J.*, 843(2):112, 2017. doi: 10.3847/1538-4357/aa7569.

[26] M. G. Aartsen et al. Differential limit on the extremely-high-energy cosmic neutrino flux in the presence of astrophysical background from nine years of IceCube data. *Phys. Rev.*, D98(6):062003, 2018. doi: 10.1103/PhysRevD.98.062003.

[27] Rafael Alves Batista, Rogerio M. de Almeida, Bruno Lago, and Kumiko Kotera. Cosmogenic photon and neutrino fluxes in the Auger era. *JCAP*, 1901(01):002, 2019. doi: 10.1088/1475-7516/2019/01/002.

[28] Alwyn Wootten and A. Richard Thompson. The Atacama Large Millimeter/submillimeter Array. *IEEE Proc.*, 97:1463, 2009. doi: 10.1109/JPROC.2009.2020572.

[29] K. I. Kellermann, R. Sramek, M. Schmidt, D. B. Shaffer, and R. Green. VLA observations of objects in the Palomar Bright Quasar Survey. *Astron. J.*, 98:1195–1207, 1989. doi: 10.1086/115207.

[30] N. Gehrels et al. The Swift Gamma-Ray Burst Mission. *Astrophys. J.*, 611:1005–1020, 2004. doi: 10.1086/422091,10.1086/427409. [Erratum: Astrophys. J.621,558(2005)].

[31] W. B. Atwood et al. The Large Area Telescope on the Fermi Gamma-ray Space Telescope Mission. *Astrophys. J.*, 697:1071–1102, 2009. doi: 10.1088/0004-637X/697/2/1071.

[32] Charles Meegan et al. The Fermi Gamma-Ray Burst Monitor. *Astrophys. J.*, 702:791–804, 2009. doi: 10.1088/0004-637X/702/1/791.

[33] H. Abdalla et al. The H.E.S.S. Galactic plane survey. *Astron. Astrophys.*, 612:A1, 2018. doi: 10.1051/0004-6361/201732098.

[34] Jamie Holder et al. The first VERITAS telescope. *Astropart. Phys.*, 25:391–401, 2006. doi: 10.1016/j.astropartphys.2006.04.002.

[35] J. Aleksic et al. Performance of the MAGIC stereo system obtained with Crab Nebula data. *Astropart. Phys.*, 35:435–448, 2012. doi: 10.1016/j.astropartphys.2011.11.007.

[36] V. F. Hess. Über Beobachtungen der durchdringenden Strahlung bei sieben Freiballonfahrten. *Phys. Z.*, 13:1084–1091, 1912.

[37] A. Achterberg et al. First Year Performance of The IceCube Neutrino Telescope. *Astropart. Phys.*, 26:155–173, 2006. doi: 10.1016/j.astropartphys.2006.06.007.

[38] A. Einstein. Über Gravitationswellen. *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften (Berlin), Seite 154-167.*, 1918.

[39] B. P. Abbott et al. Observation of Gravitational Waves from a Binary Black Hole Merger. *Phys. Rev. Lett.*, 116(6):061102, 2016. doi: 10.1103/PhysRevLett.116.061102.

[40] J. Aasi et al. Advanced LIGO. *Class. Quant. Grav.*, 32:074001, 2015. doi: 10.1088/0264-9381/32/7/074001.

[41] F. Acernese et al. Advanced Virgo: a second-generation interferometric gravitational wave detector. *Class. Quant. Grav.*, 32(2):024001, 2015. doi: 10.1088/0264-9381/32/2/024001.

[42] B. P. Abbott et al. Gravitational Waves and Gamma-rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A. *Astrophys. J.*, 848(2):L13, 2017. doi: 10.3847/2041-8213/aa920c.

[43] B. P. Abbott et al. Multi-messenger Observations of a Binary Neutron Star Merger. *Astrophys. J.*, 848(2):L12, 2017. doi: 10.3847/2041-8213/aa91c9.

[44] B. P. Abbott et al. LIGO: The Laser interferometer gravitational-wave observatory. *Rept. Prog. Phys.*, 72:076901, 2009. doi: 10.1088/0034-4885/72/7/076901.

[45] M. G. Aartsen et al. The IceCube Neutrino Observatory: Instrumentation and Online Systems. *JINST*, 12(03):P03012, 2017. doi: 10.1088/1748-0221/12/03/P03012.

[46] B. P. Abbott et al. GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Phys. Rev. Lett.*, 119(16):161101, 2017. doi: 10.1103/PhysRevLett.119.161101.

[47] B. P. Abbott et al. GW170817: Measurements of neutron star radii and equation of state. *Phys. Rev. Lett.*, 121(16):161101, 2018. doi: 10.1103/PhysRevLett.121.161101.

[48] M. G. Aartsen et al. Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A. *Science*, 361(6398):eaat1378, 2018. doi: 10.1126/science.aat1378.

[49] S. Ansoldi et al. The blazar TXS 0506+056 associated with a high-energy neutrino: insights into extragalactic jets and cosmic ray acceleration. *Astrophys. J. Lett.*, 2018. doi: 10.3847/2041-8213/aad083. [Astrophys. J.863,L10(2018)].

[50] M. G. Aartsen et al. Neutrino emission from the direction of the blazar TXS 0506+056 prior to the IceCube-170922A alert. *Science*, 361(6398):147–151, 2018. doi: 10.1126/science.aat2890.

[51] Xavier Rodrigues, Shan Gao, Anatoli Fedynitch, Andrea Palladino, and Walter Winter. Leptohadronic Blazar Models Applied to the 2014–2015 Flare of TXS 0506+056. *Astrophys. J.*, 874(2):L29, 2019. doi: 10.3847/2041-8213/ab1267.

[52] Anita Reimer, Markus Boettcher, and Sara Buson. Cascading Constraints from Neutrino Emitting Blazars: The case of TXS 0506+056. 12 2018. doi: 10.3847/1538-4357/ab2bff.

[53] K. Hirata et al. Observation of a Neutrino Burst from the Supernova SN 1987a. *Phys. Rev. Lett.*, 58:1490–1493, 1987. doi: 10.1103/PhysRevLett.58.1490. [,727(1987)].

[54] R. M. Bionta et al. Observation of a Neutrino Burst in Coincidence with Supernova SN 1987a in the Large Magellanic Cloud. *Phys. Rev. Lett.*, 58:1494, 1987. doi: 10.1103/PhysRevLett.58.1494.

[55] E. N. Alekseev, L. N. Alekseeva, I. V. Krivosheina, and V. I. Volchenko. Detection of the Neutrino Signal From SN1987A in the LMC Using the Inr Baksan Underground Scintillation Telescope. *Phys. Lett.*, B205:209–214, 1988. doi: 10.1016/0370-2693(88)91651-6.

[56] Pietro Antonioli et al. SNEWS: The Supernova Early Warning System. *New J. Phys.*, 6: 114, 2004. doi: 10.1088/1367-2630/6/1/114.

[57] S. D. Barthelmy, P. Butterworth, T. L. Cline, N. Gehrels, G. J. Fishman, C. Kouveliotou, and C. A. Meegan. BACODINE, the Real-Time BATSE Gamma-Ray Burst Coordinates Distribution Network. *Astrophysics and Space Science*, 231:235–238, September 1995. doi: 10.1007/BF00658623.

[58] M. W. E. Smith et al. The Astrophysical Multimessenger Observatory Network (AMON). *Astropart. Phys.*, 45:56–70, 2013. doi: 10.1016/j.astropartphys.2013.03.003.

[59] Hugo A. Ayala Solares et al. The Astrophysical Multimessenger Observatory Network (AMON): Performance and science program. *Astropart. Phys.*, 114:68–76, 2020. doi: 10.1016/j.astropartphys.2019.06.007.

[60] Hans Peter Dembinski, Anatoli Fedynitch, Ralph Engel, Thomas Gaisser, Felix Riehn, and Todor Stanev. Data-driven model of the cosmic-ray flux and mass composition from 10 GeV to 10^{11} GeV. *PoS*, ICRC2017:533, 2018. doi: 10.22323/1.301.0533.

[61] H. P. Dembinski et al. Report on Tests and Measurements of Hadronic Interaction Properties with Air Showers. In *Ultra High Energy Cosmic Rays (UHECR 2018) Paris, France, October 8-12, 2018*, 2019.

[62] K. A. Lave, M. E. Wiedenbeck, W. R. Binns, E. R. Christian, A. C. Cummings, A. J. Davis, G. A. de Nolfo, M. H. Israel, R. A. Leske, R. A. Mewaldt, E. C. Stone, and T. T. von Rosenvinge. Galactic Cosmic-Ray Energy Spectra and Composition during the 2009-2010 Solar Minimum Period. *Astrophys. J.*, 770:117, June 2013. doi: 10.1088/0004-637X/770/2/117.

[63] Georgia A. de Nolfo et al. Observations of the Li, Be, and B isotopes and constraints on cosmic-ray propagation. *Adv. Space Res.*, 38:1558–1564, 2006. doi: 10.1016/j.asr.2006.09.008.

[64] J. J. Engelmann, P. Ferrando, A. Soutoul, P. Goret, and E. Juliusson. Charge composition and energy spectra of cosmic-ray for elements from Be to NI - Results from HEAO-3-C2. *Astron. Astrophys.*, 233:96–111, 1990.

[65] O. Adriani et al. PAMELA Measurements of Cosmic-ray Proton and Helium Spectra. *Science*, 332:69–72, 2011. doi: 10.1126/science.1199172.

[66] O. Adriani et al. Time dependence of the proton flux measured by PAMELA during the July 2006 - December 2009 solar minimum. *Astrophys. J.*, 765:91, 2013. doi: 10.1088/0004-637X/765/2/91.

[67] M. Aguilar et al. Precision Measurement of the Helium Flux in Primary Cosmic Rays of Rigidities 1.9 GV to 3 TV with the Alpha Magnetic Spectrometer on the International Space Station. *Phys. Rev. Lett.*, 115(21):211101, 2015. doi: 10.1103/PhysRevLett.115.211101.

[68] Cristina Consolandi. Precision Measurement of the Proton Flux in Primary Cosmic Rays from 1 GV to 1.8 TV with the Alpha Magnetic Spectrometer on the International Space Station. In *25th European Cosmic Ray Symposium (ECRS 2016) Turin, Italy, September 04-09, 2016*, 2016.

[69] M. Aguilar et al. Observation of the Identical Rigidity Dependence of He, C, and O Cosmic Rays at High Rigidities by the Alpha Magnetic Spectrometer on the International Space Station. *Phys. Rev. Lett.*, 119(25):251101, 2017. doi: 10.1103/PhysRevLett.119.251101.

[70] S. P. Swordy, D. Mueller, P. Meyer, J. L'Heureux, and J. M. Grunsfeld. Relative abundances of secondary and primary cosmic rays at high energies. *Astrophys. J.*, 349:625–633, February 1990. doi: 10.1086/168349.

[71] D. Mueller, S. P. Swordy, P. Meyer, J. L'Heureux, and J. M. Grunsfeld. Energy spectra and composition of primary cosmic rays. *Astrophys. J.*, 374:356–365, June 1991. doi: 10.1086/170125.

[72] Y. S. Yoon et al. Cosmic-Ray Proton and Helium Spectra from the First CREAM Flight. *Astrophys. J.*, 728:122, 2011. doi: 10.1088/0004-637X/728/2/122.

[73] P. Maestro et al. Measurements of cosmic-ray energy spectra with the 2nd CREAM flight. *Nucl. Phys. Proc. Suppl.*, 196:239–242, 2009. doi: 10.1016/j.nuclphysbps.2009.09.045.

[74] Y. S. Yoon et al. Proton and Helium Spectra from the CREAM-III Flight. *Astrophys. J.*, 839(1):5, 2017. doi: 10.3847/1538-4357/aa68e4.

[75] P. Montini and S. M. Mari. The bending of the proton plus helium flux in primary cosmic rays measured by the ARGO-YBJ experiment in the energy range from 20 TeV to 5 PeV. 2016. arXiv:1608.01389.

[76] V. V. Prosin et al. Tunka-133: Results of 3 year operation. *Nucl. Instrum. Meth.*, A756: 94–101, 2014. doi: 10.1016/j.nima.2013.09.018.

[77] M. G. Aartsen et al. The IceCube Neutrino Observatory - Contributions to ICRC 2015 Part III: Cosmic Rays. 2015. arXiv:1510.05225.

[78] W. D. Apel et al. The spectrum of high-energy cosmic rays measured with KASCADE-Grande. 2012. arXiv:1206.3834.

[79] M. Takeda et al. Extension of the cosmic ray energy spectrum beyond the predicted Greisen-Zatsepin-Kuz'min cutoff. *Phys. Rev. Lett.*, 81:1163–1166, 1998. doi: 10.1103/PhysRevLett.81.1163.

[80] R. U. Abbasi et al. First observation of the Greisen-Zatsepin-Kuzmin suppression. *Phys. Rev. Lett.*, 100:101101, 2008. doi: 10.1103/PhysRevLett.100.101101.

[81] F. Aharonian et al. First ground based measurement of atmospheric Cherenkov light from cosmic rays. *Phys. Rev.*, D75:042004, 2007. doi: 10.1103/PhysRevD.75.042004.

[82] Patricia Sanchez-Lucas. X_{max} measurements and tests of hadronic models using the surface detector of the Pierre Auger Observatory. *PoS*, ICRC2017:495, 2018. doi: 10.22323/1.301.0495. [,48(2017)].

[83] R. U. Abbasi et al. Indications of Proton-Dominated Cosmic Ray Composition above 1.6 EeV. *Phys. Rev. Lett.*, 104:161101, 2010. doi: 10.1103/PhysRevLett.104.161101.

[84] R. U. Abbasi et al. Measurement of the Flux of Ultra High Energy Cosmic Rays by the Stereo Technique. *Astropart. Phys.*, 32:53–60, 2009. doi: 10.1016/j.astropartphys.2009.06.001.

[85] Alexander Aab et al. The Pierre Auger Observatory Upgrade - Preliminary Design Report. 2016. arXiv:1604.03637.

[86] Jörg R. Hörandel. Precision measurements of cosmic rays up to the highest energies with a large radio array at the Pierre Auger Observatory. *EPJ Web Conf.*, 210:06005, 2019. doi: 10.1051/epjconf/201921006005.

[87] Eiji Kido. The TAx4 experiment. *PoS*, ICRC2017:386, 2018. doi: 10.22323/1.301.0386.

[88] Francesco Fenu. The cosmic ray energy spectrum measured using the Pierre Auger Observatory. In *Proceedings of the 35th International Cosmic Ray Conference*, page 486. POS, 2017.

[89] Charles Jui. Summary of Results from the telescope Array Experiment. *PoS*, ICRC2015:035, 2016. doi: 10.22323/1.236.0035.

[90] Tareq AbuZayyad et al. Auger-TA energy spectrum working group report. *EPJ Web Conf.*, 210:01002, 2019. doi: 10.1051/epjconf/201921001002.

[91] R. U. Abbasi et al. Indications of Intermediate-Scale Anisotropy of Cosmic Rays with Energy Greater Than 57 EeV in the Northern Sky Measured with the Surface Detector of the Telescope Array Experiment. *Astrophys. J.*, 790:L21, 2014. doi: 10.1088/2041-8205/790/2/L21.

[92] Ralph Engel, Dieter Heck, and Tanguy Pierog. Extensive air showers and hadronic interactions at high energy. *Ann. Rev. Nucl. Part. Sci.*, 61:467–489, 2011. doi: 10.1146/annurev.nucl.012809.104544.

[93] Jose Bellido. Depth of maximum of air-shower profiles at the Pierre Auger Observatory: Measurements above $10^{17.2}$ eV and Composition Implications. In *Proceedings of the 35th International Cosmic Ray Conference*, page 506. POS, 2017.

[94] R. U. Abbasi et al. Mass composition of ultrahigh-energy cosmic rays with the Telescope Array Surface Detector data. *Phys. Rev.*, D99(2):022002, 2019. doi: 10.1103/PhysRevD.99.022002.

[95] Alexey Yushkov, Jose Bellido, John Belz, Vitor de Souza, William Hanlon, Daisuke Ikeda, Pierre Sokolsky, Yoshiki Tsunesada, and Michael Unger. Depth of maximum of air-shower profiles: testing the compatibility of measurements performed at the Pierre Auger Observatory and the Telescope Array experiment. *EPJ Web Conf.*, 210:01009, 2019. doi: 10.1051/epjconf/201921001009.

[96] Y. Fukuda et al. The Super-Kamiokande detector. *Nucl. Instrum. Meth.*, A501:418–462, 2003. doi: 10.1016/S0168-9002(03)00425-X.

[97] M. Ageron et al. ANTARES: the first undersea neutrino telescope. *Nucl. Instrum. Meth.*, A656:11–38, 2011. doi: 10.1016/j.nima.2011.06.103.

[98] S. Adrian-Martinez et al. Letter of intent for KM3NeT 2.0. *J. Phys.*, G43(8):084001, 2016. doi: 10.1088/0954-3899/43/8/084001.

[99] M. G. Aartsen et al. IceCube-Gen2: A Vision for the Future of Neutrino Astronomy in Antarctica. 2014. arXiv:1412.5106.

[100] Alexander Aab et al. Improved limit to the diffuse flux of ultrahigh energy neutrinos from the Pierre Auger Observatory. *Phys. Rev.*, D91(9):092008, 2015. doi: 10.1103/PhysRevD.91.092008.

[101] Ming-Yuan Lu, Carl G Pfendner, and Andrew Shultz. Ultra-high energy neutrino search with the Askaryan Radio Array. *PoS*, ICRC2017:966, 2018. doi: 10.22323/1.301.0966.

[102] A. Anker et al. Targeting ultra-high energy neutrinos with the ARIANNA experiment. *Adv. Space Res.*, 64:2595–2609, 2019. doi: 10.1016/j.asr.2019.06.016.

[103] Jaime Álvarez Muñiz et al. The Giant Radio Array for Neutrino Detection (GRAND): Science and Design. *Sci. China Phys. Mech. Astron.*, 63(1):219501, 2020. doi: 10.1007/s11433-018-9385-7.

[104] A. V. Olinto et al. POEMMA (Probe of Extreme Multi-Messenger Astrophysics) design. 2019. arXiv:1907.06217.

[105] Markus Ahlers and Francis Halzen. Opening a New Window onto the Universe with IceCube. *Prog. Part. Nucl. Phys.*, 102:73–88, 2018. doi: 10.1016/j.ppnp.2018.05.001.

[106] M. G. Aartsen et al. Energy Reconstruction Methods in the IceCube Neutrino Telescope. *JINST*, 9:P03009, 2014. doi: 10.1088/1748-0221/9/03/P03009.

[107] Juliana Stachurska. IceCube High Energy Starting Events at 7.5 Years – New Measurements of Flux and Flavor. *EPJ Web Conf.*, 207:02005, 2019. doi: 10.1051/epjconf/201920702005.

[108] Christian Haack and Christopher Wiebusch. A measurement of the diffuse astrophysical muon neutrino flux using eight years of IceCube data. *PoS*, ICRC2017:1005, 2018. doi: 10.22323/1.301.1005.

[109] Claudio Kopper. Observation of Astrophysical Neutrinos in Six Years of IceCube Data. *PoS*, ICRC2017:981, 2018. doi: 10.22323/1.301.0981.

[110] M. G. Aartsen et al. Observation and Characterization of a Cosmic Muon Neutrino Flux from the Northern Hemisphere using six years of IceCube data. *Astrophys. J.*, 833(1):3, 2016. doi: 10.3847/0004-637X/833/1/3.

[111] Andrea Palladino and Walter Winter. A multi-component model for observed astrophysical neutrinos. *Astron. Astrophys.*, 615:A168, 2018. doi: 10.3204/PUBDB-2018-01376,10.1051/0004-6361/201832731.

[112] M. G. Aartsen et al. Lowering IceCube's Energy Threshold for Point Source Searches in the Southern Sky. *Astrophys. J.*, 824(2):L28, 2016. doi: 10.3847/2041-8205/824/2/L28.

[113] R. Abbasi et al. An absence of neutrinos associated with cosmic-ray acceleration in γ-ray bursts. *Nature*, 484:351–353, 2012. doi: 10.1038/nature11068.

[114] Roberto Aloisio. Acceleration and propagation of ultra high energy cosmic rays. *PTEP*, 2017(12):12A102, 2017. doi: 10.1093/ptep/ptx115.

[115] Pasquale Blasi. The Origin of Galactic Cosmic Rays. *Astron. Astrophys. Rev.*, 21:70, 2013. doi: 10.1007/s00159-013-0070-7.

[116] G. Giacinti, M. Kachelrieß, and D. V. Semikoz. Explaining the Spectra of Cosmic Ray Groups above the Knee by Escape from the Galaxy. *Phys. Rev.*, D90(4):041302, 2014. doi: 10.1103/PhysRevD.90.041302.

[117] Kohta Murase, Yoshiyuki Inoue, and Charles D. Dermer. Diffuse Neutrino Intensity from the Inner Jets of Active Galactic Nuclei: Impacts of External Photon Fields and the Blazar Sequence. *Phys. Rev.*, D90(2):023007, 2014. doi: 10.1103/PhysRevD.90.023007.

[118] Xavier Rodrigues, Anatoli Fedynitch, Shan Gao, Denise Boncioli, and Walter Winter. Neutrinos and Ultra-High-Energy Cosmic-Ray Nuclei from Blazars. *Astrophys. J.*, 854(1): 54, 2018. doi: 10.3847/1538-4357/aaa7ee.

[119] Damiano Caprioli. "Espresso" Acceleration of Ultra-high-energy Cosmic Rays. *Astrophys. J.*, 811(2):L38, 2015. doi: 10.1088/2041-8205/811/2/L38.

[120] Shigeo S. Kimura, Kohta Murase, and B. Theodore Zhang. Ultrahigh-energy Cosmic-ray Nuclei from Black Hole Jets: Recycling Galactic Cosmic Rays through Shear Acceleration. *Phys. Rev.*, D97(2):023026, 2018. doi: 10.1103/PhysRevD.97.023026.

[121] Eli Waxman and John N. Bahcall. High-energy neutrinos from cosmological gamma-ray burst fireballs. *Phys. Rev. Lett.*, 78:2292–2295, 1997. doi: 10.1103/PhysRevLett.78.2292.

[122] Noemie Globus, Denis Allard, Robert Mochkovitch, and Etienne Parizot. UHECR acceleration at GRB internal shocks. *Mon. Not. Roy. Astron. Soc.*, 451(1):751–790, 2015. doi: 10.1093/mnras/stv893.

[123] Noemie Globus, Denis Allard, and Etienne Parizot. A complete model of the cosmic ray spectrum and composition across the Galactic to extragalactic transition. *Phys. Rev.*, D92 (2):021302, 2015. doi: 10.1103/PhysRevD.92.021302.

[124] Daniel Biehl, Denise Boncioli, Anatoli Fedynitch, and Walter Winter. Cosmic-Ray and Neutrino Emission from Gamma-Ray Bursts with a Nuclear Cascade. *Astron. Astrophys.*, 611:A101, 2018. doi: 10.1051/0004-6361/201731337.

[125] L. A. Anchordoqui, G. E. Romero, and J. A. Combi. Heavy nuclei at the end of the cosmic ray spectrum? *Phys. Rev.*, D60:103001, 1999. doi: 10.1103/PhysRevD.60.103001.

[126] Ke Fang, Kumiko Kotera, and Angela V. Olinto. Ultrahigh Energy Cosmic Ray Nuclei from Extragalactic Pulsars and the effect of their Galactic counterparts. *JCAP*, 1303:010, 2013. doi: 10.1088/1475-7516/2013/03/010.

[127] Luis Alfredo Anchordoqui. Acceleration of ultrahigh-energy cosmic rays in starburst superwinds. *Phys. Rev.*, D97(6):063010, 2018. doi: 10.1103/PhysRevD.97.063010.

[128] Glennys R. Farrar and Tsvi Piran. Tidal disruption jets as the source of Ultra-High Energy Cosmic Rays. 2014. arXiv:1411.0704.

[129] B. Theodore Zhang, Kohta Murase, Foteini Oikonomou, and Zhuo Li. High-energy cosmic ray nuclei from tidal disruption events: Origin, survival, and implications. *Phys. Rev.*, D96(6):063007, 2017. doi: 10.1103/PhysRevD.96.063007,10.1103/PhysRevD.96.069902. [Addendum: Phys. Rev.D96,no.6,069902(2017)].

[130] Daniel Biehl, Denise Boncioli, Cecilia Lunardini, and Walter Winter. Tidally disrupted stars as a possible origin of both cosmic rays and neutrinos at the highest energies. *Sci. Rep.*, 8(1):10828, 2018. doi: 10.1038/s41598-018-29022-4.

[131] Claire Guépin, Kumiko Kotera, Enrico Barausse, Ke Fang, and Kohta Murase. Ultra-High Energy Cosmic Rays and Neutrinos from Tidal Disruptions by Massive Black Holes. *Astron. Astrophys.*, 616:A179, 2018. doi: 10.1051/0004-6361/201732392.

[132] Kohta Murase, Kunihito Ioka, Shigehiro Nagataki, and Takashi Nakamura. High Energy Neutrinos and Cosmic-Rays from Low-Luminosity Gamma-Ray Bursts? *Astrophys. J.*, 651:L5–L8, 2006. doi: 10.1086/509323.

[133] B. Theodore Zhang, Kohta Murase, Shigeo S. Kimura, Shunsaku Horiuchi, and Peter Mészáros. Low-luminosity gamma-ray bursts as the sources of ultrahigh-energy cosmic ray nuclei. *Phys. Rev.*, D97(8):083010, 2018. doi: 10.1103/PhysRevD.97.083010.

[134] Denise Boncioli, Daniel Biehl, and Walter Winter. On the common origin of cosmic rays across the ankle and diffuse neutrinos at the highest energies from low-luminosity Gamma-Ray Bursts. *Astrophys. J.*, 872(1):110, 2019. doi: 10.3847/1538-4357/aafda7.

[135] Kohta Murase and Eli Waxman. Constraining High-Energy Cosmic Neutrino Sources: Implications and Prospects. *Phys. Rev.*, D94(10):103006, 2016. doi: 10.1103/PhysRevD.94.103006.

[136] Kenneth Greisen. End to the cosmic ray spectrum? *Phys. Rev. Lett.*, 16:748–750, 1966. doi: 10.1103/PhysRevLett.16.748.

[137] G. T. Zatsepin and V. A. Kuzmin. Upper limit of the spectrum of cosmic rays. *JETP Lett.*, 4:78–80, 1966. [Pisma Zh. Eksp. Teor. Fiz.4,114(1966)].

[138] V. S. Berezinsky and S. I. Grigor'eva. A Bump in the ultrahigh-energy cosmic ray spectrum. *Astron. Astrophys.*, 199:1–12, 1988.

[139] V. Berezinsky, A. Z. Gazizov, and S. I. Grigorieva. On astrophysical solution to ultrahigh-energy cosmic rays. *Phys. Rev.*, D74:043005, 2006. doi: 10.1103/PhysRevD.74.043005.

[140] Veniamin Berezinsky, A. Z. Gazizov, and S. I. Grigorieva. Dip in UHECR spectrum as signature of proton interaction with CMB. *Phys. Lett.*, B612:147–153, 2005. doi: 10.1016/j.physletb.2005.02.058.

[141] R. Aloisio, V. Berezinsky, Pasquale Blasi, A. Gazizov, S. Grigorieva, and B. Hnatyk. A dip in the UHECR spectrum and the transition from galactic to extragalactic cosmic rays. *Astropart. Phys.*, 27:76–91, 2007. doi: 10.1016/j.astropartphys.2006.09.004.

[142] Glennys R. Farrar. Particle Physics at Ultrahigh Energies. In *Proceedings, 18th International Symposium on Very High Energy Cosmic Ray Interactions (ISVHECRI 2014): Geneva, Switzerland, August 18-22, 2014*, 2019.

[143] A. D. Supanitsky. Implications of gamma-ray observations on proton models of ultrahigh energy cosmic rays. *Phys. Rev.*, D94(6):063002, 2016. doi: 10.1103/PhysRevD.94.063002.

[144] V. Berezinsky, A. Gazizov, and O. Kalashev. Cascade photons as test of protons in UHECR. *Astropart. Phys.*, 84:52–61, 2016. doi: 10.1016/j.astropartphys.2016.08.007.

[145] Dan Hooper and Andrew M. Taylor. On The Heavy Chemical Composition of the Ultra-High Energy Cosmic Rays. *Astropart. Phys.*, 33:151–159, 2010. doi: 10.1016/j.astropartphys.2010.01.003.

[146] R. Aloisio, V. Berezinsky, and P. Blasi. Ultra high energy cosmic rays: implications of Auger data for source spectra and chemical composition. *JCAP*, 1410(10):020, 2014. doi: 10.1088/1475-7516/2014/10/020.

[147] Andrew M. Taylor, Markus Ahlers, and Dan Hooper. Indications of Negative Evolution for the Sources of the Highest Energy Cosmic Rays. *Phys. Rev.*, D92(6):063011, 2015. doi: 10.1103/PhysRevD.92.063011.

[148] Alexander Aab et al. Combined fit of spectrum and composition data as measured by the Pierre Auger Observatory. *JCAP*, 1704(04):038, 2017. doi: 10.1088/1475-7516/2018/03/E02,10.1088/1475-7516/2017/04/038. [Erratum: JCAP1803,no.03,E02(2018)].

[149] David Wittkowski. Reconstructed properties of the sources of UHECR and their dependence on the extragalactic magnetic field. *PoS*, ICRC2017:563, 2018. doi: 10.22323/1.301.0563.

[150] R. Aloisio, D. Boncioli, A di Matteo, A. F. Grillo, S. Petrera, and F. Salamida. Cosmogenic neutrinos and ultra-high energy cosmic ray models. *JCAP*, 1510(10):006, 2015. doi: 10.1088/1475-7516/2015/10/006.

[151] Andrés Romero-Wolf and Máximo Ave. Bayesian Inference Constraints on Astrophysical Production of Ultra-high Energy Cosmic Rays and Cosmogenic Neutrino Flux Predictions. *JCAP*, 1807(07):025, 2018. doi: 10.1088/1475-7516/2018/07/025.

[152] Saikat Das, Soebur Razzaque, and Nayantara Gupta. Ultrahigh energy cosmic rays and neutrinos from light nuclei composition. *Phys. Rev.*, D99(8):083015, 2019. doi: 10.1103/PhysRevD.99.083015.

[153] Rafael Alves Batista, Denise Boncioli, Armando di Matteo, and Arjen van Vliet. Secondary neutrino and gamma-ray fluxes from SimProp and CRPropa. *JCAP*, 05:006, 2019. doi: 10.1088/1475-7516/2019/05/006.

[154] Klaes Møller, Peter B. Denton, and Irene Tamborra. Cosmogenic Neutrinos Through the GRAND Lens Unveil the Nature of Cosmic Accelerators. *JCAP*, 1905(05):047, 2019. doi: 10.1088/1475-7516/2019/05/047.

[155] Arjen van Vliet, Rafael Alves Batista, and Jörg R. Hörandel. Determining the fraction of cosmic-ray protons at ultrahigh energies with cosmogenic neutrinos. *Phys. Rev.*, D100(2): 021302, 2019. doi: 10.1103/PhysRevD.100.021302.

[156] Markus Ahlers, Luis A. Anchordoqui, and Subir Sarkar. Neutrino diagnostics of ultra-high energy cosmic ray protons. *Phys. Rev.*, D79:083009, 2009. doi: 10.1103/PhysRevD.79.083009.

[157] Graciela B. Gelmini, Oleg Kalashev, and Dmitri V. Semikoz. Gamma-Ray Constraints on Maximum Cosmogenic Neutrino Fluxes and UHECR Source Evolution Models. *JCAP*, 1201:044, 2012. doi: 10.1088/1475-7516/2012/01/044.

[158] David Wittkowski and Karl-Heinz Kampert. On the flux of high-energy cosmogenic neutrinos and the influence of the extragalactic magnetic field. *Mon. Not. Roy. Astron. Soc.*, 488(1):L119–L122, 2019. doi: 10.1093/mnrasl/slz083.

[159] Marco Stein Muzio, Michael Unger, and Glennys R. Farrar. Progress towards characterizing ultrahigh energy cosmic ray sources. *Phys. Rev. D*, 100(10):103008, 2019. doi: 10.1103/PhysRevD.100.103008.

[160] Andrew M. Taylor, Markus Ahlers, and Felix A. Aharonian. The need for a local source of UHE CR nuclei. *Phys. Rev.*, D84:105007, 2011. doi: 10.1103/PhysRevD.84.105007.

[161] Roberto Aloisio and V. S. Berezinsky. Anti-GZK effect in UHECR diffusive propagation. *Astrophys. J.*, 625:249–255, 2005. doi: 10.1086/429615.

[162] Michael Unger, Glennys R. Farrar, and Luis A. Anchordoqui. Origin of the ankle in the ultrahigh energy cosmic ray spectrum, and of the extragalactic protons below it. *Phys. Rev.*, D92(12):123001, 2015. doi: 10.1103/PhysRevD.92.123001.

[163] Ke Fang and Kohta Murase. Linking High-Energy Cosmic Particles by Black Hole Jets Embedded in Large-Scale Structures. *Nature Phys.*, 14(4):396–398, 2018. doi: 10.1038/s41567-017-0025-4.

[164] M. Kachelrieß, O. Kalashev, S. Ostapchenko, and D. V. Semikoz. Minimal model for extragalactic cosmic rays and neutrinos. *Phys. Rev.*, D96(8):083006, 2017. doi: 10.1103/PhysRevD.96.083006.

[165] A. D. Supanitsky, A. Cobos, and A. Etchegoyen. Origin of the light cosmic ray component below the ankle. *Phys. Rev.*, D98(10):103016, 2018. doi: 10.1103/PhysRevD.98.103016.

[166] Enrico Fermi. On the Origin of the Cosmic Radiation. *Phys. Rev.*, 75:1169–1174, 1949. doi: 10.1103/PhysRev.75.1169.

[167] E. Fermi. Galactic Magnetic Fields and the Origin of Cosmic Radiation. *Astrophys. J.*, 119:1–6, 1954. doi: 10.1086/145789.

[168] R. A. Treumann and C. H. Jaroschek. Fundamentals of Non-relativistic Collisionless Shock Physics: V. Acceleration of Charged Particles. 2008. arXiv:0806.4046.

[169] T. Stanev. *High Energy Cosmic Rays.* Springer Praxis Books. Springer Berlin Heidelberg, 2010. ISBN 9783540851486.

[170] Frank M. Rieger, Valenti Bosch-Ramon, and Peter Duffy. Fermi acceleration in astrophysical jets. *Astrophys. Space Sci.*, 309:119–125, 2007. doi: 10.1007/s10509-007-9466-z.

[171] Abraham Achterberg, Yves A. Gallant, John G. Kirk, and Axel W. Guthmann. Particle acceleration by ultrarelativistic shocks: Theory and simulations. *Mon. Not. Roy. Astron. Soc.*, 328:393, 2001. doi: 10.1046/j.1365-8711.2001.04851.x.

[172] Philipp Baerwald, Mauricio Bustamante, and Walter Winter. UHECR escape mechanisms for protons and neutrons from GRBs, and the cosmic ray-neutrino connection. *Astrophys. J.*, 768:186, 2013. doi: 10.1088/0004-637X/768/2/186.

[173] Yutaka Ohira, Kohta Murase, and Ryo Yamazaki. Escape-limited Model of Cosmic-ray Acceleration Revisited. *Astron. Astrophys.*, 513:A17, 2010. doi: 10.1051/0004-6361/200913495.

[174] Diego F. Torres, Gustavo E. Romero, Thomas M. Dame, Jorge A. Combi, and Yousaf M. Butt. Supernova remnants and gamma-ray sources. *Phys. Rept.*, 382:303–380, 2003. doi: 10.1016/S0370-1573(03)00201-1.

[175] Pasquale Blasi. Cosmic Ray Acceleration in Supernova Remnants. In *Cosmic rays for particle and astroparticle physics. Proceedings, 12th ICATPP Conference, Como, Italy, October 7-8, 2010*, pages 493–506, 2011. doi: 10.1142/9789814329033_0061.

[176] C. D. Dermer, S. Razzaque, J. D. Finke, and A. Atoyan. Ultrahigh Energy Cosmic Rays from Black Hole Jets of Radio Galaxies. *New J. Phys.*, 11:065016, 2009. doi: 10.1088/1367-2630/11/6/065016.

[177] Eli Waxman. Cosmological gamma-ray bursts and the highest energy cosmic rays. *Phys. Rev. Lett.*, 75:386–389, 1995. doi: 10.1103/PhysRevLett.75.386.

[178] Mario Vietri. On the acceleration of ultrahigh-energy cosmic rays in gamma-ray bursts. *Astrophys. J.*, 453:883–889, 1995. doi: 10.1086/176448.

[179] Shunsaku Horiuchi, Kohta Murase, Kunihito Ioka, and Peter Meszaros. The survival of nuclei in jets associated with core-collapse supernovae and gamma-ray bursts. *Astrophys. J.*, 753:69, 2012. doi: 10.1088/0004-637X/753/1/69.

[180] Hyesung Kang, Jorg P. Rachen, and Peter L. Biermann. Contributions to the cosmic ray flux above the ankle: Clusters of galaxies. *Mon. Not. Roy. Astron. Soc.*, 286:257, 1997. doi: 10.1093/mnras/286.2.257.

[181] Dongsu Ryu, Hyesung Kang, Eric Hallman, and T. W. Jones. Cosmological shock waves and their role in the large scale structure of the universe. *Astrophys. J.*, 593:599–610, 2003. doi: 10.1086/376723.

[182] W. F. Swann. A Mechanism of Acquirement of Cosmic-Ray Energies by Electrons. *Physical Review*, 43:217–220, February 1933. doi: 10.1103/PhysRev.43.217.

[183] Pasquale Blasi, Richard I. Epstein, and Angela V. Olinto. Ultrahigh-energy cosmic rays from young neutron star winds. *Astrophys. J.*, 533:L123, 2000. doi: 10.1086/312626.

[184] Ke Fang, Kumiko Kotera, and Angela V. Olinto. Newly-born pulsars as sources of ultrahigh energy cosmic rays. *Astrophys. J.*, 750:118, 2012. doi: 10.1088/0004-637X/750/2/118.

[185] Kumiko Kotera, Elena Amato, and Pasquale Blasi. The fate of ultrahigh energy nuclei in the immediate environment of young fast-rotating pulsars. *JCAP*, 1508(08):026, 2015. doi: 10.1088/1475-7516/2015/08/026.

[186] A. M. Hillas. The Origin of Ultrahigh-Energy Cosmic Rays. *Ann. Rev. Astron. Astrophys.*, 22:425–444, 1984. doi: 10.1146/annurev.aa.22.090184.002233.

[187] Kseniya V. Ptitsyna and Sergei V. Troitsky. Physical conditions in potential sources of ultra-high-energy cosmic rays. I. Updated Hillas plot and radiation-loss constraints. *Phys. Usp.*, 53:691–701, 2010. doi: 10.3367/UFNe.0180.201007c.0723.

[188] Eli Waxman. Cosmological origin for cosmic rays above 10**19-eV. *Astrophys. J.*, 452: L1–L4, 1995. doi: 10.1086/309715.

[189] B. Peters. Primary cosmic radiation and extensive air showers. *Il Nuovo Cimento*, 22: 800–819, November 1961. doi: 10.1007/BF02783106.

[190] Asantha Cooray. Extragalactic Background Light: Measurements and Applications. 2016. arXiv:1602.03512.

[191] A. A. Penzias and R. W. Wilson. A Measurement of Excess Antenna Temperature at 4080 Mc/s. *Astrophys. J.*, 142:419–421, July 1965. doi: 10.1086/148307.

[192] M. Tanabashi et al. Review of particle physics. *Phys. Rev. D*, 98:030001, Aug 2018. doi: 10.1103/PhysRevD.98.030001. URL `https://link.aps.org/doi/10.1103/PhysRevD.98.030001`.

[193] Justin D. Finke, Soebur Razzaque, and Charles D. Dermer. Modeling the Extragalactic Background Light from Stars and Dust. *Astrophys. J.*, 712:238–249, 2010. doi: 10.1088/0004-637X/712/1/238.

[194] R. C. Gilmore, R. S. Somerville, J. R. Primack, and A. Dominguez. Semi-analytic modeling of the EBL and consequences for extragalactic gamma-ray spectra. *Mon. Not. Roy. Astron. Soc.*, 422:3189, 2012. doi: 10.1111/j.1365-2966.2012.20841.x.

[195] Stephen K. Andrews, Simon P. Driver, Luke J. Davies, Claudia d. P. Lagos, and Aaron S. G. Robotham. Modelling the cosmic spectral energy distribution and extragalactic background light over all time. *Mon. Not. Roy. Astron. Soc.*, 474(1):898–916, 2018. doi: 10.1093/mnras/stx2843.

[196] Floyd W. Stecker, M. A. Malkan, and S. T. Scully. Intergalactic photon spectra from the far ir to the uv lyman limit for $0 < Z < 6$ and the optical depth of the universe to high energy gamma-rays. *Astrophys. J.*, 648:774–783, 2006. doi: 10.1086/506188.

[197] Alberto Franceschini, Giulia Rodighiero, and Mattia Vaccari. The extragalactic optical-infrared background radiations, their time evolution and the cosmic photon-photon opacity. *Astron. Astrophys.*, 487:837, 2008. doi: 10.1051/0004-6361:200809691.

[198] A. Dominguez et al. Extragalactic Background Light Inferred from AEGIS Galaxy SED-type Fractions. *Mon. Not. Roy. Astron. Soc.*, 410:2556, 2011. doi: 10.1111/j.1365-2966.2010.17631.x.

[199] Floyd W. Stecker, Sean T. Scully, and Matthew A. Malkan. An Empirical Determination of the Intergalactic Background Light from UV to FIR Wavelengths Using FIR Deep Galaxy Surveys and the Gamma-ray Opacity of the Universe. *Astrophys. J.*, 827(1):6, 2016. doi: 10.3847/1538-4357/aad563,10.3847/0004-637X/827/1/6. [Erratum: Astrophys. J.863,no.1,112(2018)].

[200] N. Globus, D. Allard, and E. Parizot. Propagation of high-energy cosmic rays in extragalactic turbulent magnetic fields: resulting energy spectrum and composition. *Astron. Astrophys.*, 479:97, 2008. doi: 10.1051/0004-6361:20078653.

[201] R. Aloisio and Veniamin Berezinsky. Diffusive propagation of UHECR and the propagation theorem. *Astrophys. J.*, 612:900–913, 2004. doi: 10.1086/421869.

[202] H. Bethe and W. Heitler. On the Stopping of fast particles and on the creation of positive electrons. *Proc. Roy. Soc. Lond.*, A146:83–112, 1934. doi: 10.1098/rspa.1934.0140.

[203] G. R. Blumenthal. Energy loss of high-energy cosmic rays in pair-producing collisions with ambient photons. *Phys. Rev.*, D1:1596–1602, 1970. doi: 10.1103/PhysRevD.1.1596.

[204] A. Fassò, A. Ferrari, and P. R. Sala. Total giant resonance photonuclear cross sections for light nuclei. *Proc. SATIF-3, Sendai, Japan*, pages 61–74, 1997.

[205] Alberto Fassò, Alfredo Ferrari, and Paola R. Sala. Photonuclear reactions in fluka: Cross sections and interaction models. *AIP Conference Proceedings*, 769(1):1303–1306, 2005.

[206] N. Otuka et al. Towards a More Complete and Accurate Experimental Nuclear Reaction Data Library (EXFOR): International Collaboration Between Nuclear Reaction Data Centres (NRDC). *Nucl. Data Sheets*, 120:272–276, 2014. doi: 10.1016/j.nds.2014.07.065.

[207] J. L. Puget, F. W. Stecker, and J. H. Bredekamp. Photonuclear Interactions of Ultrahigh-Energy Cosmic Rays and their Astrophysical Consequences. *Astrophys. J.*, 205:638–654, 1976. doi: 10.1086/154321.

[208] A. J. Koning, S. Hilaire, and M. C. Duijvestijn. TALYS 1.0. In *Proceedings, International Conference on Nuclear Data for Science and Tecnology*, pages 211–214, 2007.

[209] Karl-Heinz Kampert, Jörg Kulbartz, Luca Maccione, Nils Nierstenhoefer, Peter Schiffer, Günter Sigl, and Arjen Rene van Vliet. CRPropa 2.0 – a Public Framework for Propagating High Energy Nuclei, Secondary Gamma Rays and Neutrinos. *Astropart. Phys.*, 42:41–51, 2013. doi: 10.1016/j.astropartphys.2012.12.001.

[210] R. Alves Batista, D. Boncioli, A. di Matteo, A. van Vliet, and D. Walz. Effects of uncertainties in simulations of extragalactic UHECR propagation, using CRPropa and SimProp. *JCAP*, 1510(10):063, 2015. doi: 10.1088/1475-7516/2015/10/063.

[211] S. Hummer, M. Ruger, F. Spanier, and W. Winter. Simplified models for photohadronic interactions in cosmic accelerators. *Astrophys. J.*, 721:630–652, 2010. doi: 10.1088/0004-637X/721/1/630.

[212] A. Mucke, Ralph Engel, J. P. Rachen, R. J. Protheroe, and Todor Stanev. SOPHIA: Monte Carlo simulations of photohadronic processes in astrophysics. *Comput. Phys. Commun.*, 124:290–314, 2000. doi: 10.1016/S0010-4655(99)00446-4.

[213] Roberto Aloisio, Denise Boncioli, Armando Di Matteo, Aurelio F. Grillo, Sergio Petrera, and Francesco Salamida. SimProp v2r4: Monte Carlo simulation code for UHECR propagation. *JCAP*, 1711(11):009, 2017. doi: 10.1088/1475-7516/2017/11/009.

[214] Rafael Alves Batista, Andrej Dundovic, Martin Erdmann, Karl-Heinz Kampert, Daniel Kuempel, Gero Müller, Guenter Sigl, Arjen van Vliet, David Walz, and Tobias Winchen. CRPropa 3 - a Public Astrophysical Simulation Framework for Propagating Extraterrestrial Ultra-High Energy Particles. *JCAP*, 1605(05):038, 2016. doi: 10.1088/1475-7516/2016/05/038.

[215] Leonel Morejon, Anatoli Fedynitch, Denise Boncioli, Daniel Biehl, and Walter Winter. Improved photomeson model for interactions of cosmic ray nuclei. *JCAP*, 11:007, 2019. doi: 10.1088/1475-7516/2019/11/007.

[216] T. K. Gaisser and A. M. Hillas. Reliability of the method of constant intensity cuts for reconstructing the average development of vertical showers. *International Cosmic Ray Conference*, 8:353–357, 1977.

[217] Alexander Aab et al. Measurement of the average shape of longitudinal profiles of cosmic-ray air showers at the Pierre Auger Observatory. *JCAP*, 1903(03):018, 2019. doi: 10.1088/1475-7516/2019/03/018.

[218] R. U. Abbasi et al. Study of muons from ultrahigh energy cosmic ray air showers measured with the Telescope Array experiment. *Phys. Rev.*, D98(2):022002, 2018. doi: 10.1103/PhysRevD.98.022002.

[219] Sarah Müller. Direct Measurement of the Muon Density in Air Showers with the Pierre Auger Observatory. *EPJ Web Conf.*, 210:02013, 2019. doi: 10.1051/epjconf/201921002013.

[220] T. Pierog, Iu. Karpenko, J. M. Katzy, E. Yatsenko, and K. Werner. EPOS LHC: Test of collective hadronization with data measured at the CERN Large Hadron Collider. *Phys. Rev.*, C92(3):034906, 2015. doi: 10.1103/PhysRevC.92.034906.

[221] Felix Riehn, Ralph Engel, Anatoli Fedynitch, Thomas K. Gaisser, and Todor Stanev. A new version of the event generator Sibyll. *PoS*, ICRC2015:558, 2016. doi: 10.22323/1.236.0558.

[222] Sergey Ostapchenko. Monte Carlo treatment of hadronic interactions in enhanced Pomeron scheme: I. QGSJET-II model. *Phys. Rev.*, D83:014018, 2011. doi: 10.1103/PhysRevD.83.014018.

[223] D. Heck, J. Knapp, J. N. Capdevielle, G. Schatz, and T. Thouw. CORSIKA: A Monte Carlo code to simulate extensive air showers. 1998.

[224] Philipp Baerwald, Mauricio Bustamante, and Walter Winter. Are gamma-ray bursts the sources of ultra-high energy cosmic rays? *Astropart. Phys.*, 62:66–91, 2015. doi: 10.1016/j.astropartphys.2014.07.007.

[225] Eric Jones, Travis Oliphant, Pearu Peterson, et al. SciPy: Open source scientific tools for Python, 2001–. URL `http://www.scipy.org/`. [Online; accessed 2019/06].

[226] Anatoli Fedynitch, Ralph Engel, Thomas K. Gaisser, Felix Riehn, and Todor Stanev. Calculation of conventional and prompt lepton fluxes at very high energy. *EPJ Web Conf.*, 99:08001, 2015. doi: 10.1051/epjconf/20159908001.

[227] Arjen van Vliet, Jörg R. Hörandel, and Rafael Alves Batista. Cosmogenic gamma-rays and neutrinos constrain UHECR source models. *PoS*, ICRC2017:562, 2018. doi: 10.22323/1.301.0562. [35,562(2017)].

[228] Intel math kernel library developer reference. `https://software.intel.com/en-us/articles/mkl-reference-manual`. Accessed: 2019-08-09.

[229] Cuda toolkit documentation. `https://docs.nvidia.com/cuda/index.html`. Accessed: 2019-08-09.

[230] Shan Gao, Martin Pohl, and Walter Winter. On the direct correlation between gamma-rays and PeV neutrinos from blazars. *Astrophys. J.*, 843(2):109, 2017. doi: 10.3847/1538-4357/aa7754.

[231] Veniamin Berezinsky and A. Z. Gazizov. Diffusion of cosmic rays in expanding universe. *Astrophys. J.*, 643:8–13, 2006. doi: 10.1086/502626.

[232] Veniamin Berezinsky and A. Z. Gazizov. Diffusion of Cosmic Rays in the Expanding Universe. 2. Energy Spectra of Ultra-High Energy Cosmic Rays. *Astrophys. J.*, 669:684–691, 2007. doi: 10.1086/520498.

[233] Silvia Mollerach and Esteban Roulet. Magnetic diffusion effects on the ultra-high energy cosmic ray spectrum and composition. *JCAP*, 1310:013, 2013. doi: 10.1088/1475-7516/2013/10/013.

[234] F. James and M. Roos. Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations. *Comput. Phys. Commun.*, 10:343–367, 1975. doi: 10.1016/0010-4655(75)90039-9.

[235] P. Allison et al. Performance of two Askaryan Radio Array stations and first results in the search for ultrahigh energy neutrinos. *Phys. Rev.*, D93(8):082003, 2016. doi: 10.1103/PhysRevD.93.082003.

[236] Christopher Robert Persichilli. *Performance and Simulation of the ARIANNA Pilot Array, with Implications for Future Ultra-high Energy Neutrino Astronomy.* PhD thesis, University of California, Irvine, 2018. URL `https://arianna.ps.uci.edu/sites/default/files/Persichilli_Thesis.pdf`.

[237] John Krizmanic. Poemma: Probe of multi-messenger astrophysics. UHECR, 2018. URL `https://indico.in2p3.fr/event/17063/contributions/64250/`.

[238] Gunther Hasinger, Takamitsu Miyaji, and Maarten Schmidt. Luminosity-dependent evolution of soft x-ray selected AGN: New Chandra and XMM-Newton surveys. *Astron. Astrophys.*, 441:417–434, 2005. doi: 10.1051/0004-6361:20042134.

[239] Todor Stanev. Ultra high energy cosmic rays and neutrinos after Auger. 2008. arXiv:0808.1045.

[240] D. Wanderman and T. Piran. The luminosity function and the rate of Swift's gamma-ray bursts. *"Mon. Not. Roy. Astron. Soc."*, 406:1944–1958, August 2010. doi: 10.1111/j.1365-2966.2010.16787.x.

[241] Andrew M. Hopkins and John F. Beacom. On the normalisation of the cosmic star formation history. *Astrophys. J.*, 651:142–154, 2006. doi: 10.1086/506610.

[242] Cecilia Lunardini and Walter Winter. High Energy Neutrinos from the Tidal Disruption of Stars. *Phys. Rev.*, D95(12):123001, 2017. doi: 10.1103/PhysRevD.95.123001.

[243] Tobias Winchen and Stijn Buitink. Energy Spectrum of Fast Second Order Fermi Accelerators as Sources of Ultra-High-Energy Cosmic Rays. *Astropart. Phys.*, 102:25–31, 2018. doi: 10.1016/j.astropartphys.2018.04.004.

[244] Björn Eichmann, J. P. Rachen, L. Merten, A. van Vliet, and J. Becker Tjus. Ultra-High-Energy Cosmic Rays from Radio Galaxies. *JCAP*, 1802(02):036, 2018. doi: 10.1088/1475-7516/2018/02/036.

[245] Tsvi Piran. The physics of gamma-ray bursts. *Rev. Mod. Phys.*, 76:1143–1210, 2004. doi: 10.1103/RevModPhys.76.1143.

[246] A. Levan. *Gamma-Ray Bursts.* December 2018. doi: 10.1088/2514-3433/aae164.

[247] D. Band et al. BATSE observations of gamma-ray burst spectra. 1. Spectral diversity. *Astrophys. J.*, 413:281–292, 1993. doi: 10.1086/172995.

[248] Michael S. Briggs, William S. Paciesas, Geoffrey N. Pendleton, Charles A. Meegan, Gerald J. Fishman, John M. Horack, Martin Brock, Chryssa Kouveliotou, Dieter H. Hartmann, and Jon Hakkila. Batse observations of the large scale isotropy of gamma-ray bursts. *Astrophys. J.*, 459:40, 1996. doi: 10.1086/176867.

[249] Chryssa Kouveliotou, Charles A. Meegan, Gerald J. Fishman, Narayana P. Bhyat, Michael S. Briggs, Thomas M. Koshut, William S. Paciesas, and Geoffrey N. Pendleton. Identification of two classes of gamma-ray bursts. *Astrophys. J.*, 413:L101–104, 1993. doi: 10.1086/186969.

[250] Jens Hjorth and Joshua S. Bloom. The Gamma-Ray Burst - Supernova Connection. *CAPS*, 51:169–190, 2012.

[251] Edo Berger. Short-Duration Gamma-Ray Bursts. *Ann. Rev. Astron. Astrophys.*, 52:43–105, 2014. doi: 10.1146/annurev-astro-081913-035926.

[252] P. Narayana Bhat et al. The Third Fermi gbm Gamma-ray Burst Catalog: the First six Years. *Astrophys. J. Suppl.*, 223(2):28, 2016. doi: 10.3847/0067-0049/223/2/28.

[253] M. Ahlers, M. C. Gonzalez-Garcia, and F. Halzen. GRBs on probation: testing the UHE CR paradigm with IceCube. *Astropart. Phys.*, 35:87–94, 2011. doi: 10.1016/j.astropartphys.2011.05.008.

[254] Mauricio Bustamante, Philipp Baerwald, Kohta Murase, and Walter Winter. Neutrino and cosmic-ray emission from multiple internal shocks in gamma-ray bursts. 2014. doi: 10.1038/ncomms7783. [Nature Commun.6,6783(2015)].

[255] Kohta Murase, Kunihito Ioka, Shigehiro Nagataki, and Takashi Nakamura. High-energy cosmic-ray nuclei from high- and low-luminosity gamma-ray bursts and implications for multi-messenger astronomy. *Phys. Rev.*, D78:023005, 2008. doi: 10.1103/PhysRevD.78.023005.

[256] Nicholas Senno, Kohta Murase, and Peter Meszaros. Choked Jets and Low-Luminosity Gamma-Ray Bursts as Hidden Neutrino Sources. *Phys. Rev.*, D93(8):083003, 2016. doi: 10.1103/PhysRevD.93.083003.

[257] W. S. Paciesas et al. The Fourth batse gamma-ray burst catalog (revised). *Astrophys. J. Suppl.*, 122:465–495, 1999. doi: 10.1086/313224.

[258] Bing Zhang, Yi Zhong Fan, Jaroslaw Dyks, Shiho Kobayashi, Peter Meszaros, David N. Burrows, John A. Nousek, and Neil Gehrels. Physical processes shaping GRB x-ray afterglow lightcurves: Theoretical implications from the SWIFT XRT observations. *Astrophys. J.*, 642:354–370, 2006. doi: 10.1086/500723.

[259] Asaf Pe'er. Physics of Gamma-Ray Bursts Prompt Emission. *Adv. Astron.*, 2015:907321, 2015. doi: 10.1155/2015/907321.

[260] M. J. Rees and P. Meszaros. Relativistic fireballs - energy conversion and time - scales. *Mon. Not. Roy. Astron. Soc.*, 258:41–43, 1992.

[261] M. J. Rees and P. Meszaros. Unsteady outflow models for cosmological gamma-ray bursts. *Astrophys. J.*, 430:L93–L96, 1994. doi: 10.1086/187446.

[262] T. Piran. Gamma-ray bursts and the fireball model. *Phys. Rept.*, 314:575–667, 1999. doi: 10.1016/S0370-1573(98)00127-6.

[263] D. Guetta, M. Spada, and E. Waxman. Efficiency and spectrum of internal gamma-ray burst shocks. *Astrophys. J.*, 557:399, 2001. doi: 10.1086/321543.

[264] Paz Beniamini, Lara Nava, Rodolfo Barniol Duran, and Tsvi Piran. Energies of GRB blast waves and prompt efficiencies as implied by modelling of X-ray and GeV afterglows. *Mon. Not. Roy. Astron. Soc.*, 454(1):1073–1085, 2015. doi: 10.1093/mnras/stv2033.

[265] M. J. Rees and Peter Meszaros. Dissipative photosphere models of gamma-ray bursts and x-ray flashes. *Astrophys. J.*, 628:847–852, 2005. doi: 10.1086/430818.

[266] Asaf Pe'er, Peter Meszaros, and Martin J. Rees. The observable effects of a photospheric component on grb's and xrf's prompt emission spectrum. *Astrophys. J.*, 642:995–1003, 2006. doi: 10.1086/501424.

[267] R. Hascoet, F. Daigne, and R. Mochkovitch. Prompt thermal emission in gamma-ray bursts. *Astron. Astrophys.*, 551:A124, 2013. doi: 10.1051/0004-6361/201220023.

[268] D. Lazzati, B. J. Morsony, R. Margutti, and M. C. Begelman. Photospheric emission as the dominant radiation mechanism in long-duration gamma-ray bursts. *Astrophys. J.*, 765: 103, 2013. doi: 10.1088/0004-637X/765/2/103.

[269] P. Meszaros and M. J. Rees. Poynting jets from black holes and cosmological gamma-ray bursts. *Astrophys. J.*, 482:L29–L32, 1997. doi: 10.1086/310692.

[270] Bing Zhang and Huirong Yan. The Internal-Collision-Induced Magnetic Reconnection and Turbulence (ICMART) Model of Gamma-Ray Bursts. *Astrophys. J.*, 726:90, 2011. doi: 10.1088/0004-637X/726/2/90.

[271] Z. Bosnjak and P. Kumar. Magnetic jet model for GRBs and the delayed arrival of >100 MeV photons. *Mon. Not. Roy. Astron. Soc.*, 421:39, 2012. doi: 10.1111/j.1745-3933.2011.01202.x.

[272] Bo Zhang and Bing Zhang. Gamma-Ray Burst Prompt Emission Light Curves and Power Density Spectra in the ICMART Model. *Astrophys. J.*, 782:92, 2014. doi: 10.1088/0004-637X/782/2/92.

[273] He Gao and Bing Zhang. Photosphere emission from a hybrid relativistic outflow with arbitrary dimensionless entropy and magnetization in GRBs. *Astrophys. J.*, 801(2):103, 2015. doi: 10.1088/0004-637X/801/2/103.

[274] Paz Beniamini and Jonathan Granot. Properties of GRB light curves from magnetic reconnection. *Mon. Not. Roy. Astron. Soc.*, 459(4):3635–3658, 2016. doi: 10.1093/mnras/stw895.

[275] Dafne Guetta, D. Hooper, J. Alvarez-Muniz, F. Halzen, and E. Reuveni. Neutrinos from individual gamma-ray bursts in the BATSE catalog. *Astropart. Phys.*, 20:429–455, 2004. doi: 10.1016/S0927-6505(03)00211-1.

[276] Svenja Hummer, Philipp Baerwald, and Walter Winter. Neutrino Emission from Gamma-Ray Burst Fireballs, Revised. *Phys. Rev. Lett.*, 108:231101, 2012. doi: 10.1103/PhysRevLett.108.231101.

[277] Hao-Ning He, Ruo-Yu Liu, Xiang-Yu Wang, Shigehiro Nagataki, Kohta Murase, and Zi-Gao Dai. Icecube non-detection of GRBs: Constraints on the fireball properties. *Astrophys. J.*, 752:29, 2012. doi: 10.1088/0004-637X/752/1/29.

[278] Irene Tamborra and Shin'ichiro Ando. Diffuse emission of high-energy neutrinos from gamma-ray burst fireballs. *JCAP*, 1509(09):036, 2015. doi: 10.1088/1475-7516/2015/9/036,10.1088/1475-7516/2015/09/036.

[279] Eli Waxman. Gamma-ray bursts: The Underlying model. *Lect. Notes Phys.*, 598:393, 2003. arXiv:astro-ph/0303517.

[280] Francis Halzen. Lectures on neutrino astronomy: Theory and experiment. In *Neutrinos in physics and astrophysics from 10**(-33) to 10**28 CM. Proceedings, Conference, TASI'98, Boulder, USA, June 1-26, 1998*, pages 524–569, 1998.

[281] Shiho Kobayashi, Tsvi Piran, and Re'em Sari. Can internal shocks produce the variability in GRBs? *Astrophys. J.*, 490:92–98, 1997. doi: 10.1086/512791.

[282] F. Daigne and R. Mochkovitch. Gamma-ray bursts from internal shocks in a relativistic wind: temporal and spectral properties. *Mon. Not. Roy. Astron. Soc.*, 296:275, 1998. doi: 10.1046/j.1365-8711.1998.01305.x.

[283] Andrei M. Beloborodov. On the efficiency of internal shocks in gamma-ray bursts. *Astrophys. J.*, 539:L25–L28, 2000. doi: 10.1086/312830.

[284] Shiho Kobayashi and Re'em Sari. Ultra efficient internal shocks. *Astrophys. J.*, 551:934, 2001. doi: 10.1086/320249.

[285] Frederic Daigne and Robert Mochkovitch. The expected photospheric emission of GRBs in the internal shock model. *ASP Conf. Ser.*, 312:415, 2004.

[286] Z. Bosnjak, F. Daigne, and G. Dubus. Prompt high-energy emission from gamma-ray bursts in the internal shock model. *Astron. Astrophys.*, 498:677, 2009. doi: 10.1051/0004-6361/200811375.

[287] A. Mignone, G. Bodo, S. Massaglia, Titos Matsakos, O. Tesileanu, and C. Zanni. PLUTO: A Numerical Code for Computational Astrophysics. *Astrophys. J. Suppl.*, 170:228, 2007. doi: 10.1086/513316.

[288] Motoki Kino, Akira Mizuta, and Shoichi Yamada. Hydrodynamical effects in internal shock of relativistic outflows. *Astrophys. J.*, 611:1021–1032, 2004. doi: 10.1086/422305.

[289] E. Troja et al. The X-ray counterpart to the gravitational wave event GW 170817. *Nature*, 551:71–74, 2017. doi: 10.1038/nature24290. [Nature551,71(2017)].

[290] Raffaella Margutti et al. The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/VIRGO GW170817. V. Rising X-ray Emission from an Off-Axis Jet. *Astrophys. J.*, 848(2):L20, 2017. doi: 10.3847/2041-8213/aa9057.

[291] G. Hallinan et al. A Radio Counterpart to a Neutron Star Merger. *Science*, 358:1579, 2017. doi: 10.1126/science.aap9855.

[292] K. D. Alexander et al. The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/VIRGO GW170817. VI. Radio Constraints on a Relativistic Jet and Predictions for Late-Time Emission from the Kilonova Ejecta. *Astrophys. J.*, 848(2):L21, 2017. doi: 10.3847/2041-8213/aa905d.

[293] S. J. Smartt et al. A kilonova as the electromagnetic counterpart to a gravitational-wave source. *Nature*, 551(7678):75–79, 2017. doi: 10.1038/nature24303.

[294] M. Nicholl et al. The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/VIRGO GW170817. III. Optical and UV Spectra of a Blue Kilonova From Fast Polar Ejecta. *Astrophys. J.*, 848(2):L18, 2017. doi: 10.3847/2041-8213/aa9029.

[295] P. S. Cowperthwaite et al. The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/Virgo GW170817. II. UV, Optical, and Near-infrared Light Curves and Comparison to Kilonova Models. *Astrophys. J.*, 848(2):L17, 2017. doi: 10.3847/2041-8213/aa8fc7.

[296] Daniel Kasen, Brian Metzger, Jennifer Barnes, Eliot Quataert, and Enrico Ramirez-Ruiz. Origin of the heavy elements in binary neutron-star mergers from a gravitational wave event. *Nature*, 2017. doi: 10.1038/nature24453. [Nature551,80(2017)].

[297] R. Chornock et al. The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/VIRGO GW170817. IV. Detection of Near-infrared Signatures of r-process Nucle-

osynthesis with Gemini-South. *Astrophys. J.*, 848(2):L19, 2017. doi: 10.3847/2041-8213/aa905c.

[298] R. Margutti et al. The Binary Neutron Star Event LIGO/Virgo GW170817 160 Days after Merger: Synchrotron Emission across the Electromagnetic Spectrum. *Astrophys. J.*, 856 (1):L18, 2018. doi: 10.3847/2041-8213/aab2ad.

[299] E. Troja, H. van Eerten, G. Ryan, R. Ricci, J.M. Burgess, M.H. Wieringa, L. Piro, S.B. Cenko, and T. Sakamoto. A year in the life of GW 170817: the rise and fall of a structured jet from a binary neutron star merger. *Mon. Not. Roy. Astron. Soc.*, 489(2):1919–1926, 2019. doi: 10.1093/mnras/stz2248.

[300] A. Goldstein et al. An Ordinary Short Gamma-Ray Burst with Extraordinary Implications: Fermi-GBM Detection of GRB 170817A. *Astrophys. J.*, 848(2):L14, 2017. doi: 10.3847/2041-8213/aa8f41.

[301] Jens Hjorth, Andrew J. Levan, Nial R. Tanvir, Joe D. Lyman, Rados Wojtak, Sophie L. Schrøder, Ilya Mandel, Christa Gall, and Sofie H. Bruun. The Distance to NGC 4993: The Host Galaxy of the Gravitational-wave Event GW170817. *Astrophys. J.*, 848(2):L31, 2017. doi: 10.3847/2041-8213/aa9110.

[302] Jonathan Granot, Ramandeep Gill, Dafne Guetta, and Fabio De Colle. Off-Axis Emission of Short GRB Jets from Double Neutron Star Mergers and GRB 170817A. *Mon. Not. Roy. Astron. Soc.*, 481(2):1597–1608, 2018. doi: 10.1093/mnras/sty2308.

[303] Om S. Salafia, Gabriele Ghisellini, Giancarlo Ghirlanda, and Monica Colpi. Interpreting GRB170817A as a giant flare from a jet-less double neutron-star merger. *Astron. Astrophys.*, 619:A18, 2018. doi: 10.1051/0004-6361/201732259.

[304] A. Albert et al. Search for High-energy Neutrinos from Binary Neutron Star Merger GW170817 with ANTARES, IceCube, and the Pierre Auger Observatory. *Astrophys. J.*, 850(2):L35, 2017. doi: 10.3847/2041-8213/aa9aed.

[305] Shigeo S. Kimura, Kohta Murase, Peter Mészáros, and Kenta Kiuchi. High-Energy Neutrino Emission from Short Gamma-Ray Bursts: Prospects for Coincident Detection with Gravitational Waves. *Astrophys. J.*, 848(1):L4, 2017. doi: 10.3847/2041-8213/aa8d14.

[306] Ke Fang and Brian D. Metzger. High-Energy Neutrinos from Millisecond Magnetars formed from the Merger of Binary Neutron Stars. *Astrophys. J.*, 849(2):153, 2017. doi: 10.3847/1538-4357/aa8b6a. [Astrophys. J.849,153(2017)].

[307] Peter B. Denton and Irene Tamborra. Exploring the Properties of Choked Gamma-ray Bursts with IceCube's High-energy Neutrinos. *Astrophys. J.*, 855(1):37, 2018. doi: 10.3847/1538-4357/aaab4a.

[308] K. P. Mooley et al. A mildly relativistic wide-angle outflow in the neutron star merger GW170817. *Nature*, 554:207, 2018. doi: 10.1038/nature25452.

[309] Kunihito Ioka and Takashi Nakamura. Can an Off-axis Gamma-Ray Burst Jet in GW170817 Explain All the Electromagnetic Counterparts? *PTEP*, 2018(4):043E02, 2018. doi: 10.1093/ptep/pty036.

[310] Om S. Salafia, Gabriele Ghisellini, Alessio Pescalli, Giancarlo Ghirlanda, and Francesco Nappo. Light curves and spectra from off-axis gamma-ray bursts. *Mon. Not. Roy. Astron. Soc.*, 461(4):3607–3619, 2016. doi: 10.1093/mnras/stw1549.

[311] Markus Ahlers and Lea Halser. Neutrino Fluence from Gamma-Ray Bursts: Off-Axis View of Structured Jets. *Mon. Not. Roy. Astron. Soc.*, 490(4):4935–4943, 2019. doi: 10.1093/mnras/stz2980.

[312] K. P. Mooley et al. A Strong Jet Signature in the Late-time Light Curve of GW170817. *Astrophys. J.*, 868(1):L11, 2018. doi: 10.3847/2041-8213/aaeda7.

[313] E. Troja, L. Piro, G. Ryan, H. van Eerten, R. Ricci, M. Wieringa, S. Lotti, T. Sakamoto, and S. B. Cenko. The outflow structure of GW170817 from late-time broad-band observations. *Mon. Not. Roy. Astron. Soc.*, 478(1):L18–L23, 2018. doi: 10.1093/mnrasl/sly061.

[314] E. Troja et al. A luminous blue kilonova and an off-axis jet from a compact binary merger at z=0.1341. *Nature Commun.*, 9:4089, 2018. doi: 10.1038/s41467-018-06558-7.

[315] K. D. Alexander et al. A Decline in the X-ray through Radio Emission from GW170817 Continues to Support an Off-Axis Structured Jet. *Astrophys. J.*, 863(2):L18, 2018. doi: 10.3847/2041-8213/aad637.

[316] K. P. Mooley, A. T. Deller, O. Gottlieb, E. Nakar, G. Hallinan, S. Bourke, D. A. Frail, A. Horesh, A. Corsi, and K. Hotokezaka. Superluminal motion of a relativistic jet in the neutron-star merger GW170817. *Nature*, 561(7723):355–359, 2018. doi: 10.1038/s41586-018-0486-3.

[317] V. S. Ptuskin, S. I. Rogovaya, and V. N. Zirakashvili. Inverse problem for extragalactic transport of ultra-high energy cosmic rays. *JCAP*, 1503:054, 2015. doi: 10.1088/1475-7516/2015/03/054.

[318] Francesca Capel and Daniel J. Mortlock. Impact of using the ultrahigh-energy cosmic ray arrival energies to constrain source associations. *Mon. Not. Roy. Astron. Soc.*, 484(2): 2324–2340, 2019. doi: 10.1093/mnras/stz081.

[319] M. Actis et al. Design concepts for the Cherenkov Telescope Array CTA: An advanced facility for ground-based high-energy gamma-ray astronomy. *Exper. Astron.*, 32:193–316, 2011. doi: 10.1007/s10686-011-9247-0.

[320] Inés Valiño. The flux of ultra-high energy cosmic rays after ten years of operation of the Pierre Auger Observatory. In *Proceedings of the 34th International Cosmic Ray Conference*, page 271. POS, 2015.

[321] Alessio Porcelli. Measurements of Xmax above 10^{17} eV with the fluorescence detector of the Pierre Auger Observatory. In *Proceedings of the 34th International Cosmic Ray Conference*, page 420. POS, 2015.

Appendix A

Details on the PriNCe code

This appendix gives additional detail on the methods used for the PRINCE code, which is described in Chapter 4. The analytic derivations given here were also contained in the short description of PRINCE in Heinze et al. [3, app A]. The treatment and notation for the nuclear cascade follow the conventions of NEUCOSMA, which are described in Refs. [18, 124]. However, they were reformulated in a discretized way for the treatment in PRINCE.

A.1 Transport equation in detail

The transport equation for UHECRs was introduced in Chapter 3 in the simple form as:

$$\partial_t Y_i(E) = -\partial_E \left(\frac{dE}{dt} Y_i \right) - \Gamma_i Y_i(E) + \sum_j Q_{j \to i}(Y_j(E)) + J_i(E) \,. \tag{A.1}$$

We will explain the different terms in more detail in this section. The last term $J_i(E)$ simply accounts for the injection from sources for each species as a function of energy.

Continuous losses: The first term accounts for continuous energy losses. The two processes of this form are adiabatic cooling (see Section 3.2.3) and electron-positron-pair production (see Section 3.2.4), which are given by:

$$\left(\frac{dE}{dt}(z) \right)_{\mathrm{ad}} = -E H(z) \tag{A.2}$$

$$\left(\frac{dE}{dt}(z) \right)_{e^+e^-} = -\alpha r_0^2 Z^2 \, m_e^2 \int_2^\infty d\xi \, n_\gamma \left(\frac{\xi m_e}{2\Gamma}, z \right) \frac{\Phi(\xi)}{\xi^2} \,. \tag{A.3}$$

Where the latter has been derived in Blumenthal [203]. These terms account for energy losses for a single species. For the numerical solver they are treated by finite differences as explained in Section A.4.

Discrete interaction terms Discrete interactions are described by the interaction of particles and the subsequent re-injection into secondaries at different energies. Note that this can in principle also include re-injection of the same species or at the same energy (or both in the case of elastic scattering). These processes are described by the interaction rate $\Gamma_i(t)$ and the re-injection terms $Q_{j\to i}(Y_j(E))$. The reinjection terms are a short notation for the integral over the densities of all parent species:

$$Q_{j\to i}(Y_j(E)) = \int_{E_i}^{\infty} dE_j \frac{d\Gamma_{j\to i}}{dE_i}(E_j, E_i)\, Y_j(E_j) \tag{A.4}$$

Where we introduced the differential interaction rate $d\Gamma_{j\to i}/dE_i(E_j, E_i)$.

The interaction rate is given by an integral over the cross section σ_i and the target photon field n_γ:

$$\Gamma_i(E_i) = \int \mathrm{d}\varepsilon \int_{-1}^{+1} \frac{\mathrm{d}\cos\theta}{2}(1-\cos\theta) n_\gamma(\varepsilon, \cos\theta)\, \sigma_i(\varepsilon_r)\,. \tag{A.5}$$

and the inclusive interaction rate is given in the same form with the inclusive cross section $d\sigma_{i\to j}/dE_i(E_i, E_j)$:

$$\frac{d\Gamma_{j\to i}}{dE_i}(E_j, E_i) = \int \mathrm{d}\varepsilon \int_{-1}^{+1} \frac{\mathrm{d}\cos\theta}{2}(1-\cos\theta) n_\gamma(\varepsilon, \cos\theta)\, \frac{\mathrm{d}\sigma_{j\to i}}{\mathrm{d}E_i}(E_j, E_i, \varepsilon_r)\,. \tag{A.6}$$

Note that for relativistic primaries the distribution of secondaries does only depend on the ratio of energies $x = E_i/E_j$. For the implementation it is therefore practical to tabulate the inclusive cross section $d\sigma_{i\to j}/dx(x, \varepsilon_r)$ as a function of x. We omit this substitution here for clarity.

For isotropic photon fields the pitch-angle-averaged cross section $f(y)$ can be pre-computed and the interaction rate becomes:

$$\begin{aligned} \Gamma_i(E_i) &= \int \mathrm{d}\varepsilon\, n_\gamma(\varepsilon)\, f_i(y(E_i, \varepsilon)) \\ f_i(y) &= \frac{1}{2y} \int_0^{2y} \mathrm{d}\varepsilon_r\, \varepsilon_r\, \sigma_i(\varepsilon_r). \end{aligned} \tag{A.7}$$

where $y \equiv (E_i\varepsilon)/m_i$ and corresponds to the pitch-angle-averaged photon energy and the rest frame photon energy is $\varepsilon_r = y(1-\cos\theta)$. The inclusive differential cross section can again be pitch-angle-averaged and expressed as a function of y. In analogy to Eq. (A.7) the re-injection rate reads

$$\begin{aligned} Q_{j\to i}(Y_j, E_i) &= \int_{E_i}^{\infty} \mathrm{d}E_j\, Y_j(E_j, z) \int \mathrm{d}\varepsilon\, n_\gamma(\varepsilon)\, h_{j\to i}(E_i, E_j, y(E_j, \varepsilon)) \\ h_{j\to i}(E_i, E_j, y) &= \frac{1}{2y} \int_0^{2y} \mathrm{d}\varepsilon_r\ \varepsilon_r \frac{\mathrm{d}\sigma_{j\to i}}{\mathrm{d}E_i}(E_j, E_i, \varepsilon_r)\,. \end{aligned} \tag{A.8}$$

Particle decays The decay of unstable particles is governed by a term $\partial_t Y_i = -\Gamma_{\mathrm{dec},i}(E_i) Y_i$ with the decay rate $\Gamma_{\mathrm{dec},i}(E_i) = (E_i/m_i\ \tau_i)^{-1}$, where τ_i is the lifetime of an unstable particle or nucleus i at rest. The re-injection terms for the decay products have a similar form to Eq. (A.8), but do not depend on the photon field. Hence the second integral over photon energy does not appear:

$$Q_{\mathrm{dec},j\to i}(Y_j, E_i) = \int_{E_i}^{\infty} \mathrm{d}E_j\, \Gamma_{\mathrm{dec},j}(E_j)\, \frac{\mathrm{d}n_{j\to i}}{\mathrm{d}E_i}(E_j, E_i)\, Y_j(E_j)\,. \tag{A.9}$$

The redistribution function $\mathrm{d}n_{j\to i}/\mathrm{d}E_i$ is in this case the inclusive energy distributions of the decay product i in decays of j. To obtain inclusive distributions, all decay channels that contain i are summed with their branching ratio as weight.

Most unstable particles that occur in UHECR propagation have a mean lifetime much smaller than the other relevant timescales. Hence, the decay can be regarded as an instant process at the production vertex. A decay chain via the intermediate meson or nucleus u, $j \to u \to i$, can be integrated out:

$$\frac{\mathrm{d}n_{j\to u\to i}}{\mathrm{d}E_i}(E_j, E_i) = \int_{E_i}^{E_j} dE_u \frac{\mathrm{d}n_{j\to u}}{\mathrm{d}E_u}(E_j, E_u)\, \frac{\mathrm{d}n_{u\to i}}{\mathrm{d}E_i}(E_u, E_i)\,. \tag{A.10}$$

For decay chains that proceed via multiple intermediate particles this formula is applied recursively. In practice, we substitute $\mathrm{d}\sigma_{j\to u}/\mathrm{d}E_i$ (production term for the unstable particle u) in Eq. (A.8) with distributions of the decay products of u, $\mathrm{d}\sigma_{j\to u\to X}/\mathrm{d}E_X$ if $\tau_u < \tau_{\mathrm{thresh}}$. For UHECR propagation we set τ_{thresh} to ∞, i.e. all unstable particles decay immediately.

A.2 Boost conservation

A special case for the re-inejction arises for secondary nuclei. At high energies ($E_i \gg$ TeV), the impact of the internal nucleon motion can be neglected to a good approximation, resulting in the conservation of the boost of secondary fragments, i.e. the energy per nucleon is conserved.

The redistribution function then simplifies to

$$\frac{\mathrm{d}\sigma_{j\to i}}{\mathrm{d}E_i}(E_j, E_i) \approx \sigma_j\, M_{j\to i}\, \delta\left(E_i - \frac{A_i}{A_j}E_j\right), \tag{A.11}$$

where $M_{j\to i}$ is the averge multiplicity. For this case, it is convenient to express all equations in terms of energy per nucleon $E_i^A = E_i/A_i$. This leads to the simpler form of Eq. (A.11):

$$\frac{\mathrm{d}\sigma_{j\to i}}{\mathrm{d}E_i^A}(E_j^A, E_i^A) \approx \sigma_j\, M_{j\to i}\, \delta(E_i^A - E_j^A), \tag{A.12}$$

By computing the integral over this δ-distribution for the reinjection rate in Eq. (A.8) it simplifies to:

$$\begin{aligned} Q_{j\to i}(E_i^A) &= \frac{A_j}{A_i} \int \mathrm{d}\varepsilon\, n_\gamma(\varepsilon)\, Y_j(E_i^A, z)\, g_{j\to i}(y) \\ g_{j\to i}(y) &= \frac{1}{2y} \int_0^{2y} d\varepsilon_r\, \varepsilon_r\, M_{j\to i}\, \sigma(\varepsilon_r) \end{aligned} \tag{A.13}$$

For the numerical treatment, the whole system is solved by discretizing in energy, as described in Section 4.2. Therefore it is convenient to formulate the equation system in E_i^A. This makes the treatment of the δ-function in Eq. (A.12) accurate as long as the same grid in E_i^A is chosen for all nuclear particle species. We use the form Eq. (A.13) for all nuclear species in the code. However for the sake of brevity we do not mention this substitution explicitly in the formulas in Chapter 4.

A.3 Batch matrix computation of interaction rates

As mentioned in Chapter 4, the computation of the interaction and re-injection rates can be significantly accelerated by ordering them accordingly to the sparse matrix format. We will give details in this section.

By discretizing the photon field analogously to the UHECR grid, the formula for the interaction rate in Eq. (A.7) becomes:

$$\begin{aligned}\Gamma_i^k &= \Gamma_i(E_i^k)\\ &= \sum_l \Delta\varepsilon^m \, f(E_i^k, \varepsilon^m) \, n(\varepsilon^m)\\ &= \sum_l \Delta\varepsilon^m \, f_i^{km} \, n^m\\ &= \mathcal{F}_i^k \cdot \vec{n} \,.\end{aligned} \tag{A.14}$$

The factor $\Delta\varepsilon^m$ can be absorbed into either the kernel matrix $\mathcal{F}$ or the photon field vector $\vec{n}$. Here, we adopted the convention $\mathcal{F}_i^{km} = f_i^{km}\Delta\varepsilon^m$. The re-injection term from Eq. (A.8) becomes

$$\begin{aligned}\Delta E_j^l \Delta\Gamma_{ij}^{kl} &= \Delta E_j^l \frac{d\Gamma_{j\to i}}{dE_i}(E_j^l, E_i^k)\\ &= \Delta E_j^l \sum_m \Delta\varepsilon^m \, h(E_i^k, E_j^l, y^{km}) \, n(\varepsilon^m)\\ &= \Delta E_j^l \sum_m \Delta\varepsilon^m \, h_{ij}^{klm} \, n^m\\ &= (\mathcal{H}_{ij}^{kl} \cdot \vec{n}) \,,\end{aligned} \tag{A.15}$$

where the differential elements are absorbed into the kernel matrix $\mathcal{H}_{ij}^{klm} = h_{ijk}^{klm}\Delta E_j^l \Delta\varepsilon^m$. The full interaction matrix can therefore be written as:

$$\Phi_{ij}^{kl} = \begin{cases} -(\mathcal{F}_i^k \cdot \vec{n}) + (\mathcal{H}_{ij}^{kl} \cdot \vec{n}) & \text{if } i = j \text{ and } k = l \\ (\mathcal{H}_{ij}^{kl} \cdot \vec{n}) & \text{if } i \neq j \text{ or } k \neq l \end{cases} \,. \tag{A.16}$$

The computation of elements of Φ can be done in a single matrix expression if $\mathcal{F}$ and $\mathcal{H}$ are combined into a single cross section kernel $\mathcal{K}$. By ordering $\mathcal{K}$ according to the order of the $\vec{D}$ vector of Φ, the elements of $\vec{D}$ can be modified in-place without additional memory allocations:

$$\begin{pmatrix} \dots & \Phi_{ii}^{kl} & \dots & \Phi_{ij}^{kl} & \dots \end{pmatrix}^T = \mathcal{K} \cdot \begin{pmatrix} n^0 \\ \vdots \\ n^M \end{pmatrix} = \begin{pmatrix} \vdots & & \vdots \\ (-\mathcal{F}_i^{kl0} + \mathcal{H}_{ii}^{kl0}) & \dots & (-\mathcal{F}_i^{klM} + \mathcal{H}_{ii}^{klM}) \\ \vdots & & \vdots \\ \mathcal{H}_{ij}^{kl0} & \dots & \mathcal{H}_{ij}^{klM} \\ \vdots & & \vdots \end{pmatrix} \cdot \begin{pmatrix} n^0 \\ \vdots \\ n^M \end{pmatrix} \tag{A.17}$$

This arrangement allows for very fast computation of all coefficients of Φ and hence the handling of the time/redshift dependent ODE system becomes very efficient. The cross sections can be varied by scaling or replacing elements of the kernels in $\mathcal{K}$ between runs without additional initialization overhead.

A.4 Finite differences

The partial differential part of the transport equation comes with two continuous loss terms:

$$\partial_t Y = -\partial_E (b_{\mathrm{ad}} Y_i) - \partial_E \left(b_{\mathrm{pair}} Y_i\right) \tag{A.18}$$

with the loss terms defined as $b \equiv dE/dt$.

We have already accelerated the numerical solution of photo-nuclear part of the equation system by solving the ODE in the sparse matrix form. The standard approach to include partial differential terms is to express the energy derivatives as finite differences, e.g. second order central differences:

$$f'(E_i^k) = \frac{f(E_{i+1}^k) - f(E_{i-1}^k)}{2\Delta E_i^k} + \mathcal{O}((\Delta E_i^k)^2), \tag{A.19}$$

with i indicating the energy grid index. By ignoring boundaries, the discretized differentiation operator can be written on the grid as as:

$$D_i^{k,k+1} = \frac{1}{2\Delta E_i^k} \qquad D_i^{k,k-1} = \frac{-1}{2\Delta E_i^k}. \tag{A.20}$$

This is a matrix, which has only its two off-diagonals filled:

$$D = \begin{pmatrix} 0 & \frac{1}{2\Delta E^0} & & & & 0 \\ \frac{-1}{2\Delta E^0} & 0 & \frac{1}{2\Delta E^1} & & & \\ & \frac{-1}{2\Delta E^1} & 0 & \ddots & & \\ & & \ddots & 0 & \frac{1}{2\Delta E^{K-1}} & \\ & & & \frac{-1}{2\Delta E^{K-2}} & 0 & \frac{1}{2\Delta E^K} \\ 0 & & & & \frac{-1}{2\Delta E^{K-1}} & 0 \end{pmatrix} \tag{A.21}$$

Notice that this leads to an anti-symmetric matrix, which always has purely imaginary eigenvalues. Recall that the formal solution in terms of the eigenvalues reads (following Eq. (4.9)):

$$\vec{Y}(t) = \sum_i \left(e^{\lambda_i t} - 1\right) \cdot \left(\frac{C_i}{\lambda_i}\right) \vec{\Psi}_i \tag{A.22}$$

It is clear that this form of D leads to oscillations in the solution. This is typically overcome by choosing a specific solver scheme (such as Crank-Nicholson), where the ration of time step Δt and energy-grid ΔE is fixed to suppress these oscillations. However as discussed in Chapter 4, we want to use a backward differentiation scheme to deal with the stiffness of the hadronic interactions. We found by testing that it is equally accurate and more stable for our purpose to use forward biased differences, e.g. in second order:

$$f'(E_i^k) = \frac{-1f(E_i^{k+2}) + 4f(E_i^{k+1}) - 3f(E_i^k)}{2\Delta E_i^k} + \mathcal{O}((\Delta E_i^k)^2). \tag{A.23}$$

The code allows to adjust the order of finite differences to optimize for the given problem. Currently we use 6th order finite differences. While this is probably more than necessary, we find that the impact on performance is small, as the computation time is dominated by the photo-hadronic part. For applications different from UHECR propagation we might however have to revisit this choice. If the order of the operator does not change, D_i^{kl} can be included in the sparse interaction matrix Φ from Eq. (4.6).

Appendix B

Additional material from the UHECR fit

B.1 Impact of the updated 2017 data set on the two-dimensional fit

To test the methods described in Section 5.2, we use them in this appendix to reproduce the results of the Auger Combined Fit (CF) [148]. This acts mainly as a verification of our fitting methods as well as the new propagation code PriNCe. There are however some interesting implications due to the updated 2017 data set [88, 93] compared to the older 2015 data set [320, 321]. By discussing these in the exact same framework, we establish that these are purely due to the higher statistics and unrelated to other modifications that we introduced for our fit including source evolution in Section 5.3. This reproduction was part of the study published in Heinze et al. [3].

We chose exactly the same parameters as in the main results of the CF: The source evolution parameter is fixed to $m = 0$ (flat evolution); the nuclear disintegration, the CIB and the air-shower model are fixed to PSB [207], Gilmore et. al. [194] and Epos-LHC [220], respectively. The remaining free parameters are therefore the spectral index γ, R_{max} and the nuclear fractions $\mathcal{J}_A$. We minimize over these while keeping the energy scale δ_E fixed.

The fit range of the CF starts at $5 \cdot 10^9$ GeV. We noticed that with the new data set, χ^2 is significantly affected by the small discontinuity next to the $\langle X_{\mathrm{max}} \rangle$ point at $5.5 \cdot 10^9$ GeV. This point alone adds a $\chi^2 \approx 35$ to the best fit with a total $\chi^2 \approx 102$. We therefore treat this data point as an outlier and start the fit range at $6 \cdot 10^9$ GeV. This choice does not otherwise impact the fit qualitatively.

The contours for this fit are shown in Fig. B.1 and the best fit values are summarized in Tab. B.1. For the 2015 data set we find the same qualitative result as the CF: a flat extended minimum with $\gamma < 1$ and $1 \cdot 10^9 < R_{\mathrm{max}} < 8 \cdot 10^9$ GV, and a second local minimum at $\gamma \approx 2$ and $R_{\mathrm{max}} \approx 4 \cdot 10^{10}$ GV. The differences in the exact locations of the minima with respect to the CF can be partly explained by the different methods of our propagation code. The differences between SimProp and CRPropa are on a similar level, as discussed in Aab et al. [148] and Alves Batista et al. [210]. Additional small shifts originate from the handling of the experimental

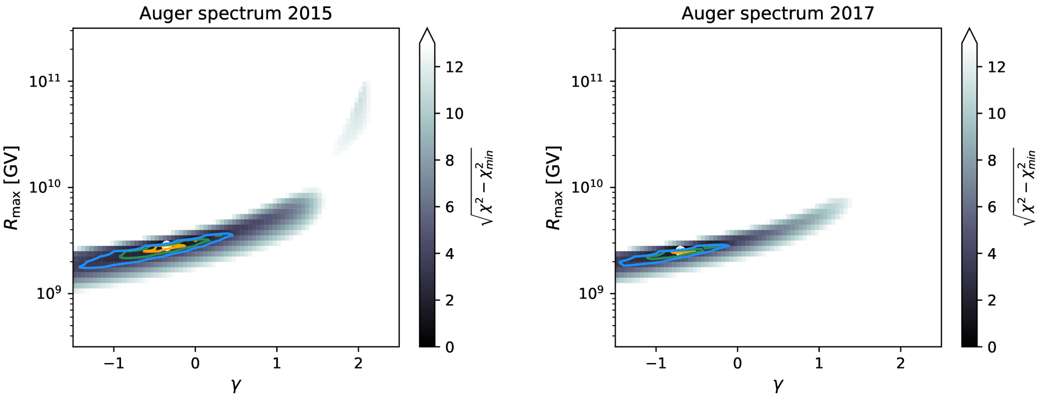

Figure B.1: Parameter space in γ and $R_{\rm max}$ assuming flat source evolution, corresponding to the Auger data sets of 2015 (left) and 2017 (right). The fits were done for model combination of PSB and Epos-LHC and with fixed energy scale, to be directly comparable to the CF baseline case [148]. The colored shading corresponds to $\sqrt{\chi^2 - \chi^2_{\rm min}}$, while $\chi^2 - \chi^2_{\rm min}$ is used to determine the contours. The contours are given for $1\sigma, 2\sigma, 3\sigma$ (for 2 d.o.f.). The best fit in each panel is indicated by a white dot. Figure taken from Heinze et al. [3]

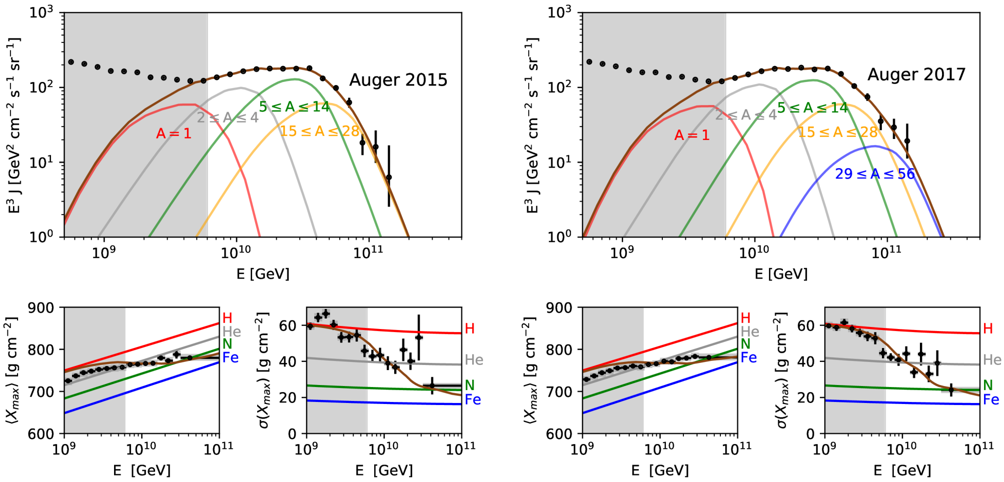

Figure B.2: Spectra (upper panels) and composition observables (lower panels) corresponding to the best fit to the Auger 2015 (left) and 2017 (right) data assuming flat source evolution and scanning in $R_{\rm max}$ and γ. The best fit values are found at $\gamma = -0.35$, $R_{\rm max} = 2.8 \cdot 10^9$ GV (2015 data) and $\gamma = -0.7$, $R_{\rm max} = 2.5 \cdot 10^9$ GV (2017 data). The grey shaded area indicates the range below $6 \cdot 10^9$ GeV, which is excluded from the fit. The expected composition is calculated assuming the EPOS-LHC shower model and comparing to the first two moments of $X_{\rm max}$ distributions. Figure taken from Heinze et al. [3]

Table B.1: Best fit parameters corresponding to the results of the fit with flat source evolution for the combination of PSB and EPOS-LHC, using the 2015 and 2017 Auger data sets. Uncertainties are given for the 1σ-interval (for 1 d.o.f.). Table taken from Heinze et al. [3]

	Auger 2015 best fit			Auger 2017 best fit		
γ	$-0.35^{+0.15}_{-0.08}$			$-0.70^{+0.12}_{-0.08}$		
R_{max} (GV)	$(2.8 \pm 0.2) \cdot 10^9$			$(2.5 \pm 0.1) \cdot 10^9$		
m	0.0 (fixed)			0.0 (fixed)		
δ_E	0.0 (fixed)			0.0 (fixed)		
$f_A(\%)$	H	He	N	H	He	N
	$5.8^{+22.0}_{-5.8}$	$89.9^{+0.6}_{-0.7}$	4.0 ± 0.2	$9.7^{+17.1}_{-9.7}$	$87.8^{+0.5}_{-0.6}$	2.4 ± 0.2
	Si	Fe		Si	Fe	
	0.3 ± 0.0	$0.0^{+4.6}_{-0.0} \cdot 10^{-3}$		0.1 ± 0.0	$(3.7 \pm 2.0) \cdot 10^{-3}$	
$I^9_A(\%)$	H	He	N	H	He	N
	$0.6^{+3.0}_{-0.6}$	$46.7^{+1.6}_{-1.8}$	$39.9^{+1.2}_{-1.3}$	$0.8^{+1.9}_{-0.8}$	$47.9^{+1.3}_{-1.4}$	$37.9^{+1.5}_{-1.6}$
	Si	Fe		Si	Fe	
	$12.8^{+1.1}_{-1.2}$	$0.0^{+1.0}_{-0.0}$		$11.4^{+2.2}_{-2.3}$	2.1 ± 1.1	
χ^2 / dof	44.4 / 22			65.3 / 22		

observables. While we fit the first two X_{max} moments for the composition, the CF uses the full X_{max} distribution. This has the strongest impact on the second minimum at $\gamma = 2$, which becomes less significant in our approach. In addition, we directly fit the unfolded spectrum and do not use a forward-folding procedure in the fit. We cannot reproduce the exact same approach without access to the Auger internal data (full X_{max} distribution, detailed detector response).

When switching to the 2017 data set, the best fit parameters do not qualitatively change (see Tab. B.1). However, the χ^2_{min} becomes worse due to the higher statistics. The allowed contours become narrower with a stronger preference for positive spectral indices. The second local minimum disappears. The reasons are the reduced statistical errors and a narrower width of the X_{max} distribution at the highest energies of the 2017 data set, leaving less room for the combination of a high R_{max} with somewhat softer spectral indices.

The largest qualitative difference concerns the injected iron fraction. While the 2015 data set does not require iron at the source, the new data suggest a small - but non-zero - integral iron fraction $I^9_{\mathrm{Fe}} \approx 2\%$. This is also visible in the comparison of the best fit spectra in Fig. B.2: for the 2017 data set (right panel) there is a contribution of heavy elements at the cutoff, which is absent in the fit to the 2017 data set (left panel). This is due to the increased statistics of the three highest energy data points in the spectrum, which have a higher flux in the updated data set. The fit is mainly driven by the fit to the *ankle*, where the data has the highest statistics.

The fit in this range fixes the lighter elements, leading to low maximal rigidities. As we chose the maximal energy proportional to the charge Z, sustaining the flux at the the highest energy data points then requires an iron fraction. However, this relies on the initial assumptions, i.e. rigidity dependence of the maximal energy and the fixed energy scale. Hence it cannot be rigorously interpreted as evidence for a non-zero iron fraction. Note however that this is also found in the three dimensional fit for all model combinations as demonstrated in Section 5.3.4. An indication for an iron contribution can also be inferred directly in the composition data above $10^{19.4}$ eV as discussed in Unger [8]. However, that result depends more strongly on the hadronic interaction model.

B.2 Additional plots and data from the 3D fit

This appendix provides additional material from the UHECR parameter fit in Chapter 5: Tab. B.2 contains the best fit parameters for the nine different model combinations. Fig. B.3 and Fig. B.4 show the fit contours in $\gamma - R_{\mathrm{max}}$ and $R_{\mathrm{max}} - m$ respectively. The best fit spectrum, composition and injection spectrum for the combination of Peanut - Epos-LHC are shown in Fig. B.5 and Fig. B.6.

Table B.2: Best fit parameters for the 3D parameter scan with free source evolution for all nine model combinations, as described in Section 5.3.2

		TALYS - SIBYLL 2.3	TALYS - EPOS-LHC	TALYS - QGSJETII-04
γ		$-0.80^{+0.27}_{-0.23}$	$-0.05^{+0.10}_{-1.45}$	$-1.40^{+0.07}_{-0.10}$
$R_{\max}$ (GV)		$(1.6 \pm 0.2) \cdot 10^9$	$2.5^{+0.0}_{-0.9} \cdot 10^9$	$1.8^{+0.2}_{-0.0} \cdot 10^9$
m		$4.2^{+0.4}_{-0.6}$	$-6.0^{+8.0}_{-0.0}$	$-6.0^{+0.2}_{-0.0}$
δ_E		$0.14^{+0.00}_{-0.03}$	$0.11^{+0.03}_{-0.01}$	$0.14^{+0.00}_{-0.01}$
$f_A(\%)$	H	$0.0^{+42.6}_{-0.0}$	$0.0^{+14.3}_{-0.0}$	$82.2^{+1.3}_{-1.5}$
	He	$82.0^{+3.8}_{-6.4}$	$90.0^{+0.4}_{-0.4}$	$17.3^{+0.9}_{-0.9}$
	N	$17.3^{+1.0}_{-1.1}$	$9.8^{+0.4}_{-0.4}$	$0.5^{+0.0}_{-0.0}$
	Si	$0.6^{+0.1}_{-0.1}$	$0.3^{+0.1}_{-0.1}$	$0.0^{+0.0}_{-0.0}$
	Fe	$0.0^{+0.0}_{-0.0}$	$0.0^{+0.0}_{-0.0}$	$0.0^{+0.0}_{-0.0}$
$I^9_A(\%)$	H	$0.0^{+1.2}_{-0.0}$	$0.0^{+1.6}_{-0.0}$	$12.3^{+1.0}_{-1.1}$
	He	$9.8^{+2.8}_{-2.9}$	$38.1^{+0.9}_{-1.0}$	$27.6^{+1.2}_{-1.3}$
	N	$69.2^{+1.5}_{-1.6}$	$54.2^{+1.0}_{-1.1}$	$58.3^{+0.5}_{-0.5}$
	Si	$17.9^{+3.2}_{-3.5}$	$6.4^{+1.8}_{-1.8}$	$0.0^{+1.4}_{-0.0}$
	Fe	$3.2^{+1.2}_{-1.3}$	$1.4^{+0.7}_{-0.7}$	$1.8^{+0.3}_{-0.3}$
χ^2 / dof		27.0 / 21	53.1 / 21	259.1 / 21

		PSB - SIBYLL 2.3	PSB - EPOS-LHC	PSB - QGSJETII-04
γ		$-1.50^{+0.55}_{-0.00}$	$0.75^{+0.12}_{-0.09}$	$-1.50^{+0.05}_{-0.00}$
$R_{\max}$ (GV)		$1.4^{+0.5}_{-0.0} \cdot 10^9$	$3.5^{+0.5}_{-0.4} \cdot 10^9$	$1.8^{+0.1}_{-0.2} \cdot 10^9$
m		$5.0^{+0.4}_{-0.6}$	$-6.0^{+0.4}_{-0.0}$	$-6.0^{+0.2}_{-0.0}$
δ_E		$0.14^{+0.00}_{-0.11}$	$0.14^{+0.00}_{-0.03}$	$0.14^{+0.00}_{-0.02}$
$f_A(\%)$	H	$0.0^{+37.2}_{-0.0}$	$0.0^{+5.9}_{-0.0}$	$83.7^{+0.7}_{-0.8}$
	He	$98.5^{+0.1}_{-0.1}$	$87.8^{+0.2}_{-0.2}$	$16.1^{+0.4}_{-0.4}$
	N	$1.4^{+0.2}_{-0.2}$	$11.1^{+0.6}_{-0.6}$	$0.2^{+0.0}_{-0.0}$
	Si	$0.1^{+0.0}_{-0.0}$	$1.0^{+0.3}_{-0.3}$	$0.0^{+0.0}_{-0.0}$
	Fe	$0.0^{+0.0}_{-0.0}$	$0.1^{+0.1}_{-0.1}$	$0.0^{+0.0}_{-0.0}$
$I^9_A(\%)$	H	$0.0^{+1.8}_{-0.0}$	$0.0^{+1.5}_{-0.0}$	$19.4^{+0.9}_{-0.9}$
	He	$34.7^{+1.3}_{-1.3}$	$55.9^{+0.5}_{-0.5}$	$42.7^{+0.8}_{-0.8}$
	N	$40.7^{+2.5}_{-2.7}$	$34.9^{+1.4}_{-1.4}$	$34.8^{+1.0}_{-1.0}$
	Si	$21.2^{+3.4}_{-3.7}$	$7.6^{+2.0}_{-2.1}$	$0.9^{+1.6}_{-0.9}$
	Fe	$3.4^{+1.7}_{-1.7}$	$1.5^{+0.9}_{-1.0}$	$2.1^{+0.8}_{-0.8}$
χ^2 / dof		23.8 / 21	46.6 / 21	228.8 / 21

		PEANUT - SIBYLL 2.3	PEANUT - EPOS-LHC	PEANUT - QGSJETII-04
γ		$-0.75^{+0.34}_{-0.21}$	$-1.50^{+0.08}_{-0.00}$	$-1.50^{+0.03}_{-0.00}$
$R_{\max}$ (GV)		$1.8^{+0.3}_{-0.1} \cdot 10^9$	$1.6^{+0.2}_{-0.0} \cdot 10^9$	$1.8^{+0.2}_{-0.0} \cdot 10^9$
m		$3.4^{+0.6}_{-0.6}$	$0.6^{+0.6}_{-0.8}$	$-6.0^{+0.2}_{-0.0}$
δ_E		$0.01^{+0.03}_{-0.04}$	$0.14^{+0.00}_{-0.01}$	$0.14^{+0.00}_{-0.00}$
$f_A(\%)$	H	$0.0^{+18.8}_{-0.0}$	$62.3^{+5.8}_{-8.3}$	$84.7^{+0.8}_{-0.9}$
	He	$93.8^{+0.5}_{-0.5}$	$37.1^{+1.2}_{-1.3}$	$15.1^{+0.5}_{-0.5}$
	N	$5.7^{+0.5}_{-0.5}$	$0.7^{+0.0}_{-0.0}$	$0.2^{+0.0}_{-0.0}$
	Si	$0.4^{+0.1}_{-0.1}$	$0.0^{+0.0}_{-0.0}$	$0.0^{+0.0}_{-0.0}$
	Fe	$0.0^{+0.0}_{-0.0}$	$0.0^{+0.0}_{-0.0}$	$0.0^{+0.0}_{-0.0}$
$I^9_A(\%)$	H	$0.0^{+0.9}_{-0.0}$	$5.2^{+1.4}_{-1.4}$	$17.2^{+0.9}_{-0.9}$
	He	$24.9^{+1.6}_{-1.6}$	$35.3^{+1.2}_{-1.3}$	$34.9^{+1.0}_{-1.0}$
	N	$47.8^{+2.2}_{-2.4}$	$50.3^{+1.3}_{-1.4}$	$44.4^{+1.0}_{-1.1}$
	Si	$24.5^{+3.0}_{-3.3}$	$8.0^{+2.2}_{-2.3}$	$2.3^{+1.5}_{-1.6}$
	Fe	$2.8^{+1.3}_{-1.4}$	$1.2^{+0.9}_{-0.9}$	$1.3^{+0.6}_{-0.7}$
χ^2 / dof		32.9 / 21	38.5 / 21	209.9 / 21

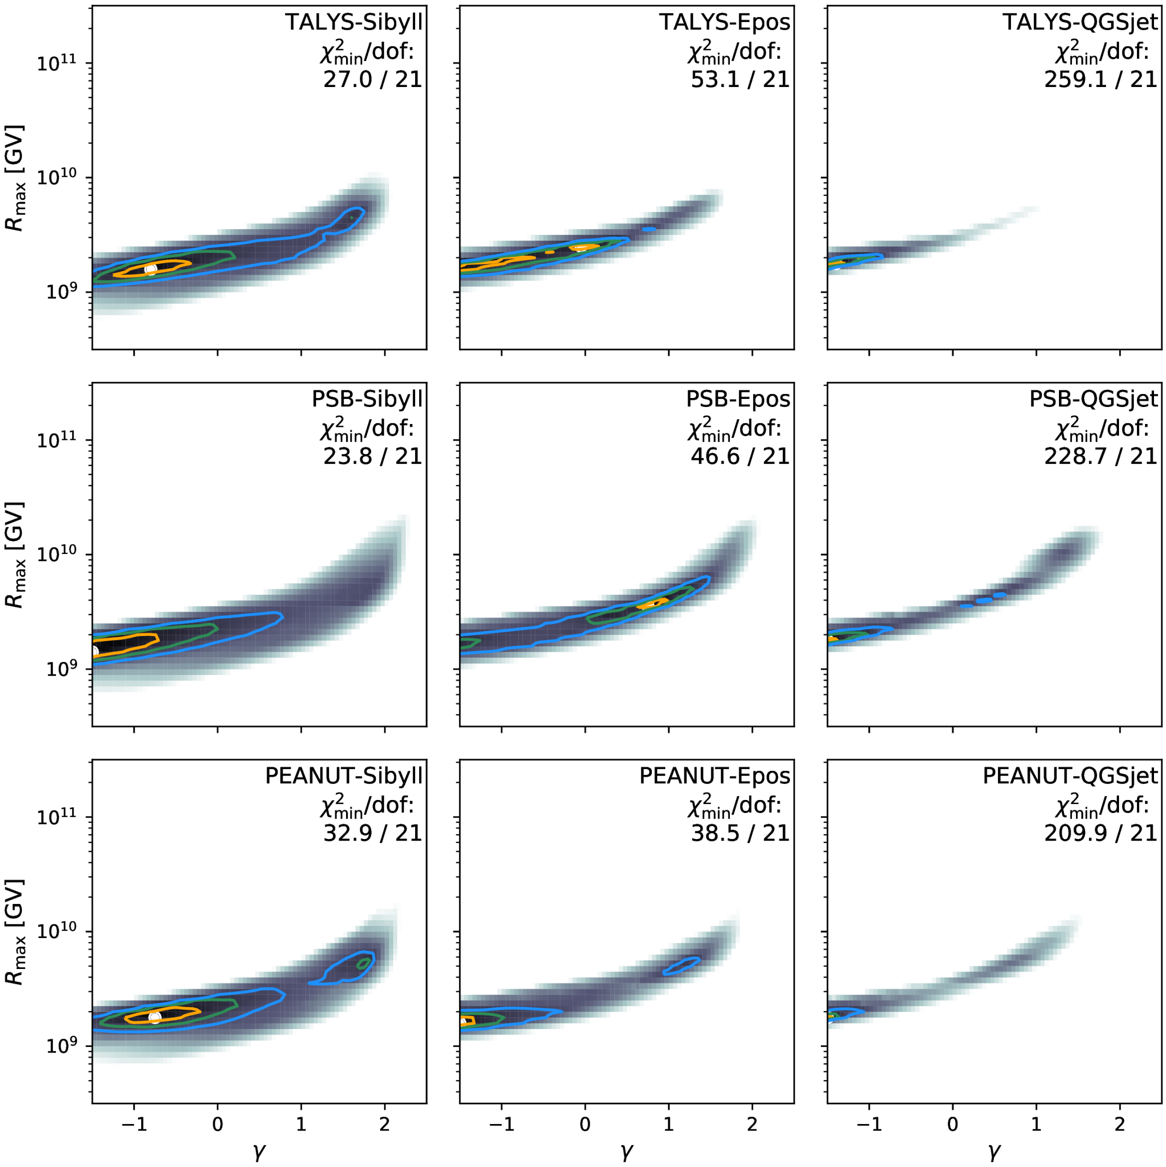

Figure B.3: Same as Fig. 5.5 but for the plane in $\gamma - R_{\rm max}$.

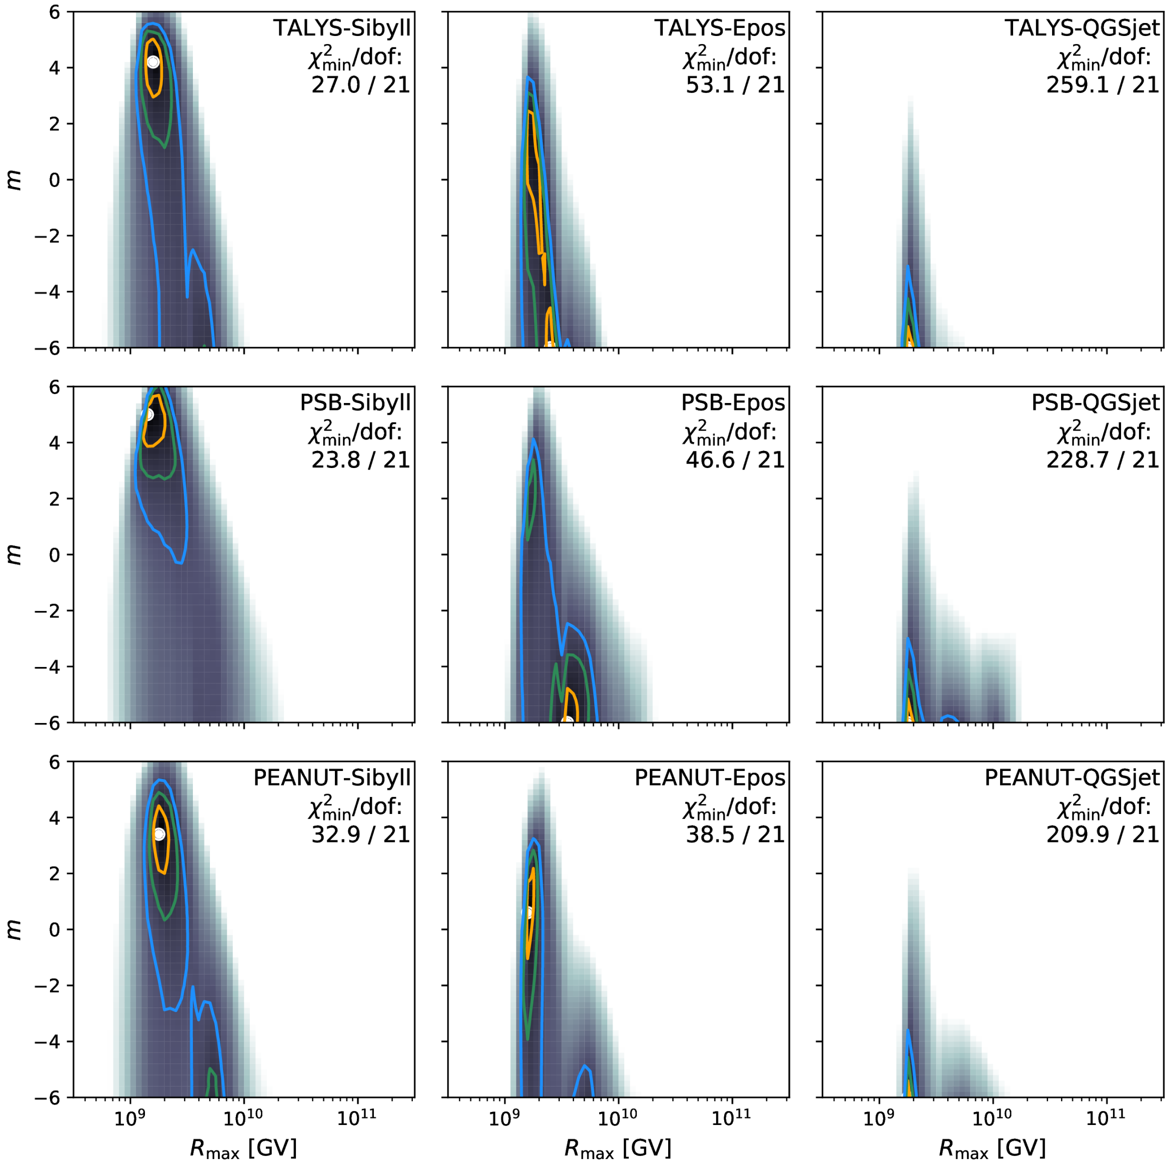

Figure B.4: Same as Fig. 5.5 but for the plane in $R_{\text{max}} - m$.

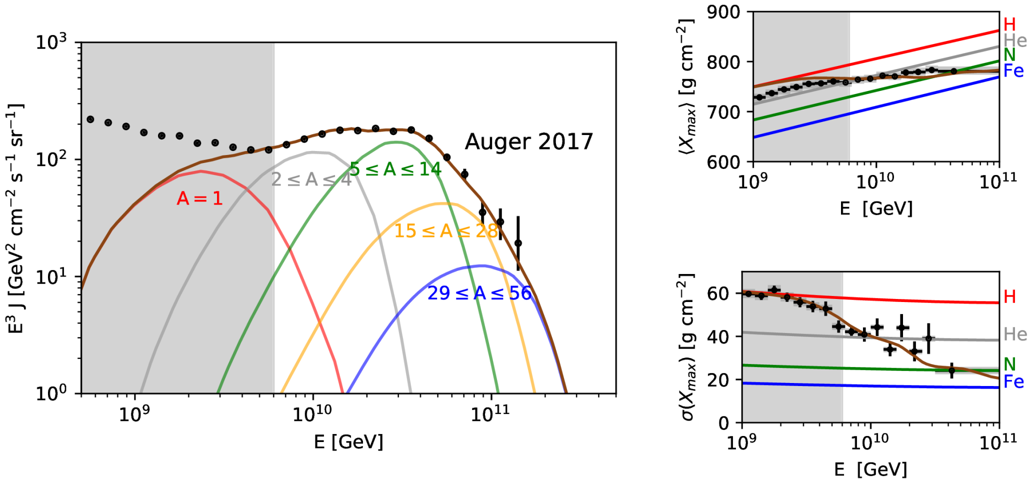

Figure B.5: Same as Fig. 5.1, but for the combination of PEANUT - EPOS-LHC

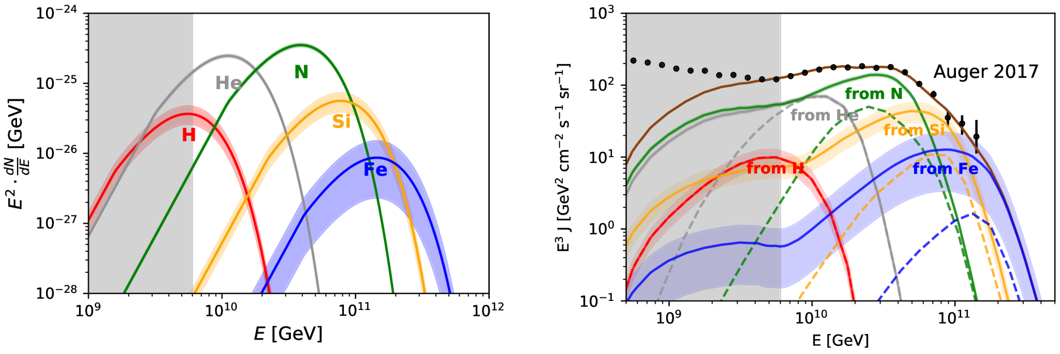

Figure B.6: Same as Fig. 5.2, but for the combination of PEANUT - EPOS-LHC