Blazars as Sources of Neutrinos and Ultra-high-energy Cosmic Rays

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. Xavier Rodrigues

Präsident 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. Dr Ana Franckowiak
3. Prof Dr Markus Böttcher

Tag der mündlichen Prüfung: 29. August 2019

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. 126/2014 am 18.11.2014 angegebenen Hilfsmittel angefertigt habe.

Xavier Rodrigues
Berlin, den 26. März 2019

List of publications

Rodrigues, X., Fedynitch, A., Gao, S., Boncioli, D., & Winter, W., *Neutrinos and Ultra-High-Energy Cosmic-Ray Nuclei from Blazars* (2018), Astrophys. J., 854 (54), arXiv:1711.02091

Palladino, A., Rodrigues, X., Gao, S., & Winter, W., *Interpretation of the diffuse astrophysical neutrino flux in terms of the blazar sequence* (2019), Astrophys. J., 871 (41), arXiv:1806.04769

Rodrigues, X., Biehl, D., Boncioli, D., & Taylor, A. M., *Binary neutron star merger remnants as sources of cosmic rays below the "Ankle"* (2019), Astropart. Phys., 106 (10), arXiv:1806.01624

Rodrigues, X., Gao, S., Fedynitch, A., Palladino, A., & Winter, W., *Leptohadronic blazar models applied to the 2014-15 flare of TXS 0506+056* (2018), accepted for publication in Astrophys. J. Lett., arXiv:1812.05939

Abstract

The origin of ultra-high-energy cosmic rays (UHECRs) is still unclear, but there is evidence that their origin lies in extragalactic sources. At the same time, neutrino telescopes like IceCube have observed a flux of high-energy astrophysical neutrinos, expected to originate in cosmic ray (CR) interactions. However, the arrival directions of the observed neutrinos do not seem to significantly correlate with the coordinates of known high-energy astrophysical sources. In this thesis we contribute to the understanding of this problem by exploring blazars, a class of active galactic nuclei (AGNs), as potential sites for the acceleration and interaction of UHECRs. Motivated by evidence that a fraction of the observed UHECRs are heavier than protons, we model numerically the interactions of a population of accelerated nuclei with the environment photon fields present in blazars. We then estimate the emitted neutrinos and UHECR spectrum and composition. We conclude that in low-luminosity blazars, accelerated CRs do not interact efficiently due to the low density of the photon fields, but instead escape the source unscathed, while in high-luminosity blazars (such as flat-spectrum radio quasars, FSQRs), photo-hadronic and photo-nuclear interactions are efficient, leading to abundant neutrino production and the development of a nuclear cascade of secondary nuclei that are lighter than the accelerated isotope. We then use our model to quantify the neutrino emission from the entire cosmological distribution of blazars. We conclude that a population of low-luminosity blazars, currently unobserved but expected theoretically, can explain the entire IceCube flux at the highest energies. However, if that is the case, then high-luminosity blazars must have a comparatively low hadronic content in order to explain the lack of correlations between neutrinos and bright gamma-ray sources. We also model neutrino and photon emission from one particular blazar, object TXS 0506+056, from whose direction a neutrino was recently detected during a state of enhanced electromagnetic activity. We test the hypothesis that a signal of 13 ± 5 muon neutrinos observed by IceCube from the same direction in 2014-15 may have originated in the same source. Given the constraints from multi-wavelength observations, we show that such photo-hadronic models can explain at most 5 events observed by IceCube, which seems to disfavor the hypothesis that this blazar was the source of the signal. Finally, we turn our attention from blazars to the remnants of neutron star mergers, and study their potential as CR emitters. The only neutron star merger ever observed was detected recently in gravitational waves, and its remnant has since been monitored by telescopes in different wavelengths. We model the non-thermal interactions in the source and show that radio and X-ray observations, as well as the non-observation of the remnant in gamma rays, can provide crucial constraints on the magnetic field strength. Given these constraints, we estimate that this source class is capable of accelerating and emitting very-high-energy CRs. This result emphasizes the importance of future gravitational wave observations to better constrain the population of these sources. While the results of this thesis provide steps towards an understanding of CR and neutrino production, further work is still necessary, including a joint source-propagation model capable of better constraining the sources based on UHECR data.

Zusammenfassung

Der Ursprung ultra-hochenergetischer kosmischer Strahlung (UHECRs) ist immer noch unbekannt, aber es gibt Hinweise, dass ihr Ursprung extragalaktische Quellen sind. Zusätzlich messen Neutrinoteleskope wie IceCube einen Fluss hochenergetischer astrophysikalischer Neutrinos, dessen erwarteter Ursprung Wechselwirkungen kosmischer Strahlung (CR) ist. Jedoch scheinen die Ankunftsrichtungen der beobachteten Neutrinos nicht signifikant mit den Koordinaten bekannter, hochenergetischer astrophysikalischer Quellen zu korrelieren. In dieser Dissertation tragen wir zum Verständnis dieses Problems durch die Untersuchung von Blazaren, eine Klasse aktiver Galaxienkerne (AGNs), bei. Dort können UHECRs potentiell beschleunigt werden sowie auch wechselwirken. Motiviert durch Hinweise, dass ein Teil der beobachteten UHECRs schwerer als Protonen ist, modellieren wir die Wechselwirkungen einer Population beschleunigter Kerne mit den umgebenden Photonfelder in Blazaren numerisch. Wir folgern, dass in Blazaren niedriger Luminosität beschleunigte CRs nicht effizient wechselwirken. Auf der anderen Seite sind photo-hadronische und photo-nukleare Wechselwirkungen effizient in Blazaren hoher Luminosität, was zu starker Neutrinoproduktion und zur Entwicklung einer nuklearen Kaskade führt, die leichter als das beschleunigte Isotop sind. Anschließend nutzen wir unser Modell um die Neutrinoemission der gesamten kosmologischen Verteilung von Blazaren zu quantifizieren. Wir folgern, dass eine Population von Blazaren niedriger Luminosität, die derzeit nicht beobachtet, aber theoretisch erwartet wird, den gesamten IceCube-Fluss bei den höchsten Energien erklären kann. Sollte dies der Fall sein, hätten Blazare hoher Luminosität einen vergleichbar niedrigen hadronischen Inhalt um die Abwesenheit von Korrelationen zwischen Neutrinos und hellen Gamma-Strahlungsquellen zu erklären. Weiterhin modellieren wir die Neutrinoemission des Blazars TXS-0506+056, aus dessen Richtung ein Neutrino während einer Phase erhöhter elektromagnetischer Aktivität detektiert wurde. Wir testen die Hypothese, dass ein Signal von 13 ± 5 Myon-Neutrinos, die in IceCube aus der selben Richtung im Jahr 2014-15 gemessen wurden, von der selben Quelle stammt. Wir zeigen, dass solche photo-hadronischen Modelle höchstens 5 Ereignisse erklären können, was die Hypothese dieses Blazars als Quelle ausschließen würde. Schließlich wenden wir uns Überresten von verschmelzenden Neutronensternen zu und untersuchen ihr Potential als CR-Emitter. Das einzige Ereignis verschmelzender Neutronensterne, das bisher beobachtet wurde, wurde kürzlich in Gravitationswellen detektiert. Der Überrest wurde danach von verschiedenen Teleskopen in verschiedenen Wellenlängen überwacht. Wir modellieren nichtthermische Wechselwirkungen in der Quelle und zeigen, dass Radio- und Röntgenmessungen, als auch die Nichtbeobachtung des Überrests in Gamma-Strahlung strikte Beschränkungen der magnetischen Feldstärke nach sich zieht. Wir zeigen, dass diese Quelle in der Lage ist, CRs auf sehr hohe Energien zu beschleunigen und zu emittieren. Dieses Ergebnis betont die Wichtigkeit von zukünftigen Messungen von Gravitationswellen, um diese Quellen besser einzuschränken. Diese Dissertation trägt zum Verständnis von CR- und Neutrinoproduktion bei, jedoch sind weitere Studien nötig, vor Allem ein kombiniertes Quelle-Propagation-Modell, das die Quellen auf der Basis von UHECR-Daten besser einschränken kann.

Acknowledgments

First and foremost I would like to thank my advisor, Walter Winter, for his supervision and support, and for teaching me about the science and craft of multi-messenger astrophysics.

A word of deep gratitude to Shan Gao, Andrea Palladino, Denise Boncioli, Anatoli Fedynitch and Arjen van Vliet for their careful reading of parts of the manuscript and invaluable comments. Any errors left are obviously my own. I also wish to thank the post-docs and students of the NEUCOS project, as well as Andrew M. Taylor, for collaborating and sharing their knowledge with me during my stay at DESY.

I want to extend my gratitude to Daniel Biehl, Chun-Sung Jao and Leonel Morejon for their companionship and far-ranging discussions over the last three years.

A very special thanks to my Dad for his love and support. The appreciation that he and my Mum always had for knowledge and culture is the reason I stand where I am today.

Finally, to my wife Ana for her unwavering support, both emotional and intellectual, and for reading parts of the manuscript. As always, this work is dedicated to her.

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 abbreviations

AGN	active galactic nucleus
BAT	Burst Alert Telescope
BL Lac	BL Lacertae
BLR	broad line region
BR	branching ratio
CC	charged current
CMB	cosmic microwave background
CR	cosmic ray
DOM	digital optical module
DSA	diffusive shock acceleration
EC	external Compton
EGRB	extragalactic gamma-ray background
FIR	far infrared
FSRQ	flat-spectrum radio quasar
GBM	gamma-ray burst monitor
GW	gravitational wave
GZK	Greisen-Zatsepin-Kuzmin
HESE	high-energy starting event
HSP	high-synchrotron peaked
IC	inverse Compton
IGMF	intergalactic magnetic field
IGRB	isotropic diffuse gamma-ray background
IR	infrared
ISP	intermediate-synchrotron peaked
KN	Klein-Nishina
LAT	Large-Area Telescope
LSP	low-synchrotron peaked
MHD	magneto-hydro-dynamic
NC	neutral current
PDE	probability distribution function
SED	spectral energy distribution
SSC	synchrotron self-Compton
UHE	ultra-high-energy
UHECR	ultra-high-energy cosmic ray
UV	ultraviolet
VHE	very-high-energy

Contents

Chapter 1

Introduction

Cosmic rays (CRs) are the most energetic particles in nature. There is evidence that at the highest energies, they consist not only of protons but also heavier nuclei. In spite of numerous CR measurements that have provided information about their spectrum and composition, the origin of this cosmic radiation is not yet fully understood. There is clear evidence that some sources within our own galaxy, like supernova remnants, pulsars and pulsar wind nebulae are capable of accelerating CRs up to $\sim$ PeV energies; on the other hand, there is now evidence that ultra-high-energy CRs (UHECRs), *i.e.* with energy above 10^{18} eV, should originate in extragalactic sources (Abbasi et al., 2017; Aab et al., 2017a). However, there if of yet no self-consistent picture of their origin, which must include an understanding of their acceleration and interactions in a distribution of sources, and their subsequent propagation to Earth.

A potential key to unlocking this conundrum is the fact that neutrinos should be produced when UHECRs interact with surrounding radiation, both inside the astrophysical source and during propagation in the Universe. Thus, neutrino emission represents the smoking-gun evidence for hadronic sources. In 2013, the IceCube neutrino telescope first detected a flux of neutrinos whose high energies strongly suggest their astrophysical origin (Aartsen et al., 2013). However, after years of accumulated data, the arrival directions of these astrophysical neutrinos are not statistically correlated with the positions in the sky of known high-energy sources (Aartsen et al., 2018a). From a source-modeling perspective this is unexpected, since the processes that lead to neutrino emission in UHECR sources are expected to co-produce high-energy radiation like X-rays and gamma rays. On the other hand, the detection of gamma ray sources is limited by the sensitivity of current telescopes, which means the observation of the high-energy Universe is confined to sources that are either close by or bright enough to be detected from high redshifts. Furthermore, both CRs and high-energy photons lose energy through interactions with the isotropic photon backgrounds; on the contrary, neutrinos can travel cosmological distances unscathed due to their small interaction cross sections, losing energy only to the adiabatic expansion of the Universe. This means that neutrinos can probe deep into the cosmos, while both CRs and very-high-energy (VHE) photons are limited to the local Universe (*cf.* Sec-

tions 2.1.2 and 2.3.1). Considering these facts, it is perhaps less surprising that the first ~ 100 astrophysical neutrinos ever observed are weakly correlated with the limited source sample.

Nevertheless, the increasing amount of data available on neutrinos and UHECRs can provide crucial constraints to the properties of their sources and on the mechanisms at play in the high-energy Universe. The source classes commonly considered to be in the origin of this cosmic radiation are powerful extragalactic objects, like blazars (Protheroe, 1995; Essey et al., 2010; Murase et al., 2014; Padovani et al., 2016), the cores of active galactic nuclei (AGNs, Stecker, 2013), gamma-ray bursts (GRBs, Paczynski & Xu, 1994; Waxman & Bahcall, 1997; Vietri, 1998; Meszaros & Waxman, 2001; Hummer et al., 2012; Murase & Ioka, 2013), or star-forming galaxies (Loeb & Waxman, 2006; Stecker, 2013; Tamborra et al., 2014). The common denominator of these sources is the capability to accelerate a thermal pool of particles to UHEs (in shocks between fast plasma outflows or through turbulent magnetic fields), and then make these particles radiate through non-thermal processes.

Throughout most of this work we will focus on blazars, a sub-class of AGNs, which are violent emissions of matter and radiation powered by accreting supermassive black holes in the center of certain galaxies. In the case of blazars, the relativistic jet launched by the black hole is pointing directly towards Earth, leading to the increase of the observed luminosity due to relativistic beaming and to the Doppler boost of the emission, whose electromagnetic spectrum spans a broad range of frequencies, ranging from radio to TeV energies. In fact, blazars are the dominant gamma-ray point sources in the sky beyond our own galaxy (Ackermann et al., 2015a), and are also commonly believed to dominate the isotropic diffuse gamma-ray background Di Mauro et al. (2014). The multi-wavelength emission from blazars is typically explained in two frameworks. In *leptonic* models, the high-energy radiation results from the up-scattering of soft photons by a non-thermal population of electrons accelerated within the source; because these models involve purely electromagnetic processes, they cannot explain either CR or neutrino emission. A second distinct framework is that of *lepto-hadronic* models, where gamma-ray emission originates from secondary particles produced in hadronic interactions of UHE protons or nuclei with soft photons in the source, leading to the production of mesons, which decay into neutrinos, electron-positron pairs and gamma-ray photons. The escape of these high-energy hadrons provides an explanation of UHECR emission, while the hadronic interactions lead to the emission of neutrinos and gamma rays by the source. Therefore, this class of models can in principle explain simultaneously the emission of high-energy photons, UHECRs and neutrinos, although the same object may not be an efficient emitter of all three types of radiation. For example, a high density of soft photons in the source may increase neutrino production by boosting hadronic interactions, while at the same time attenuating gamma ray emission through photon annihilation, turning the source opaque to its own high-energy radiation (Murase et al., 2016). This shows that the study of

neutrino production in blazars requires a self-consistent modeling of all the processes at play in the source.

Photo-hadronic blazar models have been applied extensively to explain the observed multi-wavelength emission of a number of blazars, and the expected neutrino emission can also be estimated. This is either done in analytical or semi-analytical approaches, or by means of more sophisticated numerical calculations. The numerical framework typically involves sampling the redistribution functions of the secondary particles produced in hadronic interactions using Monte Carlo codes (like SOPHIA, Mücke et al., 2000) and feed this information into a system of kinetic equations of all the particles species involved.

All the above-mentioned models deal with the radiation processes undergone by a non-thermal population of protons and electrons; however, as mentioned earlier, the high-energy tail of the CR spectrum has been suggested to have a component of nuclei heavier than protons, with an average mass between that of pure protons and pure iron-56 (Aab et al., 2017b). While some works have addressed the interactions of heavier nuclei in AGNs, they usually treat neutrino production in an overly simplified fashion. For example, Anchordoqui et al. (2008) have adopted an analytical estimation for the shape of the AGN neutrino spectrum, and used the total observed CR energy flux to normalize the total CR emissivity of the entire AGN population. However, as shown in Chapter 3, the spectrum and total luminosity of neutrinos emitted by blazars depends critically on the composition of the accelerated CRs, which must be addressed in a consistent numerical treatment of all the photo-nuclear and photo-hadronic interactions in the source. In Chapter 3 I present the results of a new numerical model of CR interactions in blazar jets, which includes nuclei heavier than protons (Rodrigues et al., 2018a). This is the first model to numerically calculate the effect of the blazar photon fields in all relevant photo-nuclear processes including neutrino production. Additionally, the thermal and molecular emission external to the jet that takes place in a category of AGNs (*cf.* Section 2.3) is taken into account regarding its effects in CR cooling and neutrino emission. This model is therefore an important tool for studying the processes behind the emission of neutrinos and CR nuclei heavier than protons from blazars.

Such model of photo-hadronic interactions in a single source can be generalized and applied to an entire population of blazars. Given the lack of statistical correlations between neutrino arrival directions and the positions of gamma-ray blazars, Aartsen et al. (2017a) have excluded the contribution of these sources to the IceCube neutrino flux at a level any higher than 20%. However, as argued earlier, this sample corresponds to only a fraction of the entire cosmological distribution of blazars, due to the limited sensitivity of gamma-ray telescopes. Motivated by these arguments, in Chapter 4 I will show how our blazar neutrino model can be used to estimate the diffuse neutrino emission from the entire blazar population. The conclusion is that under

certain conditions, a sub-set of low-luminosity blazars is able to account for the entire spectrum of astrophysical neutrinos detected by IceCube in the range from sub-PeV to PeV energies (Palladino et al., 2019). At the same time, a limit can be set to the amount of CRs accelerated in high-luminosity blazars, due to the lack of overlap between neutrinos and gamma-ray blazars.

In spite of this overall mismatch between the map of gamma-ray blazars and that of astrophysical neutrinos, there have been sporadic evidence in the past of correlations between the directions of IceCube neutrino events and some specific blazars (Resconi et al., 2017; Padovani et al., 2016; Kadler et al., 2016). Most recently, in September 2017, IceCube observed a high-energy neutrino from the direction of the blazar TXS 0506+056, which was found to be simultaneously in a state of gamma-ray flaring (Aartsen et al., 2018b). This temporal and spacial coincidence led to a high confidence level of the association between the neutrino and the gamma-ray flare. This event was therefore regarded as the first discovery of a (non-stellar) neutrino source, and became one of the major recent events in the field of astroparticle physics. There have since been several proposals for a phenomenology of this flare (*e.g.* Gao et al., 2018; Keivani et al., 2018; Cerruti et al., 2015, see Section 2.5), capable of describing the emission of a neutrino spectrum compatible with the observed event, as well as the multi-wavelength emission observed simultaneously.

In searching for further evidence of neutrino emission from TXS 0506+056, Aartsen et al. (2018c) have found an enhanced neutrino flux from the direction of the source during a six-month period in 2014–15. When combining these data with the knowledge that a coincidence was observed in 2017, the significance of a neutrino flare from TXS in 2014–15 lies at the 3.5 σ level. However, archival data from the Fermi LAT gamma-ray telescope reveals that essentially no enhancement was taking place simultaneously in the gamma-ray emission of the source (Aartsen et al., 2019), which makes the neutrino flare a major challenge from a theoretical standpoint. Chapter 5 is dedicated to this problem. There, I present several models for the hadronic interactions in TXS 0506+056 during the 2014–15 flare (Rodrigues et al., 2018b), ranging from a one-zone model to more geometrically complex models, involving different radiation zones in the jet (*cf.* Section 2.3) or including the interactions of high-energy protons with a thermal field emitted from the accretion disk of the central black hole. My particular contribution to this work was the development of the external field model. The general conclusion is that the emission of a number of neutrinos compatible with the observed flare is challenging in the photo-hadronic paradigm, since it would imply the triggering of photo-hadronic cascades in the source that would lead to the brightening of the X-ray and gamma-ray emission, which would be in conflict with the observed data. In another contemporary work, Reimer et al. (2018) have also focused on the archival neutrino flare and their results are generally in agreement with this conclusion.

A more complete picture of CR emission from blazars must combine a model of radiative

interactions in the source with propagation effects, in order to self-consistently estimate the observed CR, photon and neutrino flux from these sources. This topic is the focus of a project in progress at the time of writing of this thesis, and its results must therefore be left outside our scope.

In the final chapter, we discuss what was perhaps the most striking breakthrough in recent multi-messenger astronomy: the observation in August 2017 of a long-duration gravitational wave (GW) event (Abbott et al., 2017a) which was also identified as a short GRB (Abbott et al., 2017b; Savchenko et al., 2017). GWs are a phenomenon predicted by General Relativity, consisting of small distortions in the curvature of spacetime that propagate at the speed of light. In the astronomic context, recent observations have proved them to be a surprisingly common phenomena, closely related to the merging events of compact binary systems. The first gravitational waves had already been detected in September 2015 by the LIGO and Virgo collaborations, as the result of a merger of two black holes (Abbott et al., 2016), and since then nearly a dozen other black hole mergers have been detected. However, the 2017 event, dubbed GW170817, remains the only GW event to date to be also observed in its electromagnetic counterpart. Indeed, two seconds after the GW signal, the Fermi Gamma-ray Burst Monitor (GBM) was triggered by a gamma-ray signal (Goldstein et al., 2017). The source was identified as the first ever observed neutron star merger, in a galaxy at distance of 40 Mpc. During the following year, an extensive follow-up campaign was launched to probe the multi-wavelength emission from the remnant of the merger. A continuum of non-thermal emission was found, spanning from radio to X-rays (Hallinan et al., 2017; Alexander et al., 2017; Margutti et al., 2017; Ruan et al., 2018), as well as bright thermal emission in the optical band (Valenti et al., 2017), associated with nuclear reactions of heavy isotopes.

Like in blazars, the power-law emission is a signature of a population of accelerated particles, most likely electrons in this case, that radiate non-thermally. This shows that there is efficient particle acceleration within the remnant, making it a potential candidate for a site of CR acceleration. At the same time, the observation of a kilonova from nuclear emission confirms the presence of heavy elements within the source, which can also be co-accelerated and possibly emitted as CRs. Because of all these aspects, the last chapter of this thesis is dedicated to the modeling of this merger remnant (Rodrigues et al., 2019). The continuous non-thermal spectral emission is typically attributed to radiative processes by leptons (since electrons are much more radiative than nuclei). I have therefore developed a code for leptonic interactions for the modeling of the electromagnetic emission from the remnant. The first conclusion is that the non-thermal fluxes observed in the first ~ 100 days can be explained by synchrotron emission of a population of electrons accelerated in the source up to $\gtrsim 10$ TeV. The same electron population should then up-scatter the environment radiation and emit gamma rays. I will then show that

the non-observation of the source by the H.E.S.S. gamma-ray telescope can place a lower limit on the magnetic field strength in the remnant. Given the range of parameters obtained from the leptonic diagnosis, the source is shown to possess the conditions for accelerating CR protons up to energies of several PeV to several EeV at late times, while heavier nuclei can be accelerated to even higher energies due to their higher electric charge. This conclusion has important implications, since, as pointed out by Hillas (2004), there should be an additional component in the CR spectrum around 10^{18} eV (just below the ankle of the spectrum, *cf.* Section 2.1.1) making the bridge between the low-energy range, dominated by Galactic sources, and the UHE range. These results show that binary neutron star merger remnants are good candidates for acceleration sites of CRs at those energies. Similar results have also been obtained by another independent, simultaneous study (Kimura et al., 2018).

The remainder of this thesis is structured as follows. Chapter 2 serves a more in-depth introduction to the relevant observational data, namely multi-wavelength blazar observations, UHECRs, and neutrinos, as well as the current status of blazar radiation modeling. In Chapter 3 I present a novel lepto-hadronic blazar model including interactions of nuclei heavier than protons, and show the effect of CR composition in the emitted neutrino spectrum (Rodrigues et al., 2018a). In Chapter 4, the model presented before is applied to blazars of different luminosities in order to estimate the contribution from the entire distribution of blazars to the IceCube diffuse astrophysical neutrino flux (Palladino et al., 2019). In Chapter 5 I discuss the particular case of TXS 0506+056. Lepto-hadronic models are applied to the 2014–15 neutrino flare (Rodrigues et al., 2018b), and the challenges and shortcomings of each model are analyzed, taking into account the constraints from multi-wavelength observations. In Chapter 6 I depart from the topic of blazars and focus on the modeling of the remnant of the neutron star merger GW170817 (Rodrigues et al., 2019). This source class is shown to be a potential source of CRs below the ankle. Chapter 7 offers a conclusion to this thesis.

Chapter 2

Multi-messenger and blazar astrophysics

2.1 Ultra-high-energy cosmic ray observations

2.1.1 Spectrum and composition

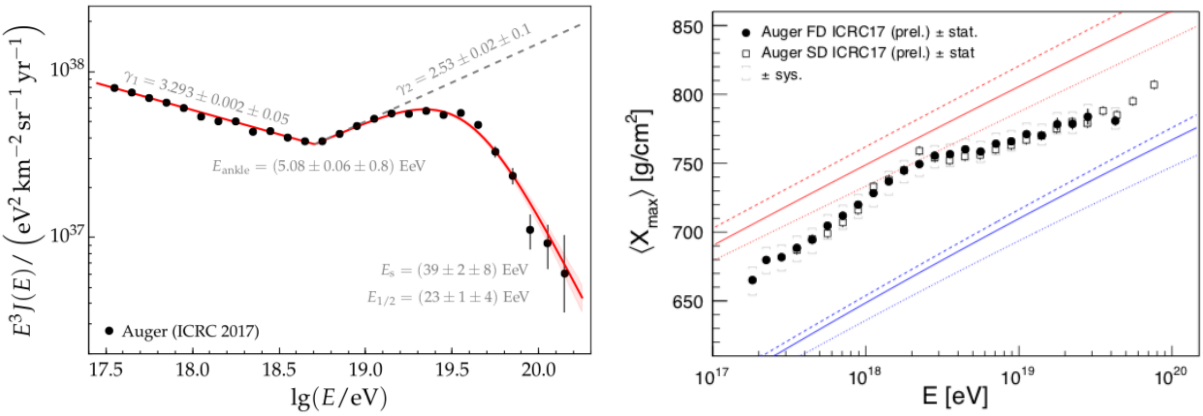

Figure 2.1: *Left:* CR spectrum above $10^{17.5}$ eV measured by the Pierre Auger Observatory. The red line represents a fit with a broken power law (represented by the dashed gray line) and a suppression at UHEs. *Right:* Average depth of the shower maximum, $\langle X_{\mathrm{max}} \rangle$, as a function of energy. For comparison the predictions of different hadronic models are shown as lines for protons (red) and iron (blue). Both plots taken from Unger (2018).

Cosmic ray observations can probe the nature of high-energy astrophysical sources where particle acceleration takes place. In fact, we know currently that CRs are accelerated up to at least 100 EeV, an extraordinary amount of energy to be carried by one subatomic particle.

The CR spectrum has now been measured by several experiments across twelve orders of magnitude in energy. Although UHECRs were first detected as early as 1962 (Linsley, 1963), only recently has the high-energy end of the spectrum been accurately measured by collecting information on extensive air showers triggered in the atmosphere. There are currently two

UHECR experiments in operation, namely the Pierre Auger Observatory (Aab et al., 2015a), in Argentina, and the Telescope Array (TA, Kawai et al., 2008), in the state of Utah, USA. A big challenge in UHECR measurements is the fact that at the highest energies CRs are exceedingly rare; at 100 EeV their flux is as low as one particle per square kilometer per century. Moreover, because the detection method is indirect, the measurements rely on hadronic interaction models, which contribute with a large systematic uncertainty (De Ridder et al., 2018; Dembinski et al., 2019). On the flip side, the data obtained from the observation of these extensive air showers can be used to test these hadronic models at energies much higher than those reached by particle accelerators (Aab et al., 2016).

The left panel of Figure 2.1 shows the spectrum of CRs above 300 PeV measured by Auger. The spectral hardening at around 5 EeV is called the *ankle*, one of the features of the CR spectrum. Another important feature not represented in the figure is a softening at around 3 PeV dubbed the *knee*. In general, the presence of these features indicates an energy of transition in the sources or mechanisms that dominate the production of CRs, or a process that affects CRs at that energy during propagation. The ankle is thought to be the point where extragalactic sources start dominating CR emission (Aloisio et al., 2012), while at lower energies the contribution from Galactic sources must dominate. It has also been proposed that the ankle results instead from energy losses of extragalactic protons due to interactions with the cosmic microwave background (CMB) during their propagation (De Marco et al., 2003; Berezinsky et al., 2006; Aloisio et al., 2007). However, Heinze et al. (2016) have shown that this explanation of the ankle is disfavored, since the flux of neutrinos expected from these interactions (named *cosmogenic* neutrinos) would be in conflict with the established upper limits (Aab et al., 2015b; Ishihara, 2016, for more recent limits see Aartsen et al. (2018d)).

Regardless of the nature of the ankle feature itself, there is now evidence that the CRs in the high-energy tail of the spectrum have indeed extragalactic origin (Abbasi et al., 2017; Aab et al., 2017a). Besides the observational evidence, this conclusion also follows from theoretical considerations, insofar as neither the Milky Way or any know source within it has the conditions to efficiently accelerate and contain such high-energy CRs, namely large enough size or strong enough magnetic fields (Hillas, 1984, *cf.* Section 2.1.2). However, the question is still open of exactly what types of extragalactic sources are responsible for the observed UHECR flux, and what mechanisms are at play in their acceleration. That question is one of the motivations for the present work.

Besides the ankle, the other feature in the UHECR spectrum is a cutoff at ultra-high energies. Recent analyses suggest that this feature should correspond to the maximum energy of the CR spectrum emitted by the astrophysical sources (Aab et al., 2017b), although in principle it may also originate in the cooling of UHECRs due to interactions with the isotropic photon

background during their propagation in the intergalactic medium (see Section 2.1.3).

Regardless of the type of source where UHECRs are accelerated, the observed abundance of UHECRs can be used to estimate the total power emitted by a local distribution of these sources, or local *emissivity*, expressed as the power in UHECRs output by an isotropic distribution of sources per unit cosmological volume. For a CR spectrum scaling with E^{-2}, the local (*i.e.* $z = 0$) emissivity of CRs above 10^{19} eV has been estimated to be $\sim 10^{44}$ erg yr^{-1}Mpc^{-3} (Waxman, 1995).

Measurements of UHECRs are indirect, and rely on the hadronic and electromagnetic air showers created by CRs when they hit the top of the atmosphere. An array of detectors spread across a large area of land detects the charged particles in these showers, and the energy of the primary CR is reconstructed. Hybrid detectors like Auger and TA combine ground-level air shower measurements with fluorescence detectors that observe the development of the air shower by detecting ultraviolet (UV) radiation from the nitrogen fluorescence caused by the charged particles of the showers. This allows for the direct measurement of the energy loss as a function of the depth in the atmosphere. In the right panel of Figure 2.1 we show the Auger measurements of the average *depth of the shower maximum*, or $\langle X_{\rm max} \rangle$, as a function of the detected energy. The value of $X_{\rm max}$ is the distance between the top of the atmosphere and the point where the air shower reaches the maximum number of secondary particles, before getting thinner again due to energy losses. On average, a heavier CR isotope will trigger the air shower higher up in the atmosphere; therefore, on average, heavier elements create showers with lower values of $X_{\rm max}$. The average value of this quantity can therefore be used as a proxy for the mass of the primary CRs at a given energy. The red and blue lines are the estimates of the $\langle X_{\rm max} \rangle$ for protons and iron, respectively, estimated by current hadronic models. We can see that the values of $\langle X_{\rm max} \rangle$ measured are not consistent with the theoretical predictions for a pure proton composition (red lines), but rather indicates a mixed composition that gets gradually heavier above 2 EeV, corresponding to an average mass between helium and iron. This suggests that there is a component of nuclei heavier than protons in the measured sample of UHECRs (Aab et al., 2017b).

Most current models of hadronic interactions in blazars, introduced in Section 2.4, are based on the acceleration of protons; however, based on Auger data, it seems to be the case that heavier nuclei are being accelerated together with protons at least in some sources, which has consequences for the photon and neutrino emission from those sources. This was the motivation behind the blazar model presented in Chapter 3, where we consider the effect of the acceleration of nuclei heavier than protons in blazars (Rodrigues et al., 2018a). Of course, a complete picture of the sources of CRs must ultimately combine a source model with the effects of intergalactic propagation. These effects are briefly introduced in Section 2.1.3, although such combined

source-propagation model lies outside the scope of this thesis.

2.1.2 Cosmic ray acceleration

The processes that accelerate CRs up to ultra-high energies (UHEs) are still a matter of intense study and debate. Diffusive shock acceleration (DSA), also called first-order Fermi acceleration, a model originally developed in the 1960s, is nowadays the most widely accepted mechanism behind the efficient acceleration of non-thermal CRs in astrophysical sources (*e.g.* Drury, 1983; Blandford & Eichler, 1987; Jones & Ellison, 1991). In the DSA model, the acceleration of a charged particle occurs in supersonic shocks in a collisionless plasma that carries a magnetic field. When the particle crosses the shock front (*i.e.* the interface between two plasma flows), it draws kinetic energy from the shock and accelerates. The magnetic field in the medium must then return the particle back to the shock front and the particle crosses it again in the opposite direction. This back-and-forth diffusive motion repeats, gradually accelerating he particle until the shock front eventually leaves it behind (see *e.g.* Baring, 1997).

The importance of the shock acceleration mechanism for particle astrophysics lies in the fact that a population of particles accelerated by this process (as well as other processes like second-order Fermi acceleration) will follow a power law spectrum (Fermi, 1949), which is a necessary feature to explain the observed CR spectrum. The spectral index of CRs from DSA depends on the compression ratio of the shock (*i.e.* the ratio between the up- and downstream flow speeds); in strong shocks the emerging spectrum has an index of $\alpha = 2$, while weaker shocks accelerate the particles to softer spectra, $\alpha > 2$.

Besides shock acceleration, other processes may contribute to CR acceleration in astrophysical sources, like stochastic acceleration in sites with turbulent magnetic fields or acceleration by magnetic reconnection (de Gouveia Dal Pino & Kowal, 2013), which occurs when two magnetic flows with opposite polarity annihilate, forming current sheets where particles can be accelerated. These processes should be particularly relevant in acceleration sites that are dominated by the magnetic field rather than matter, and generally lead to CR spectra that are harder compared to DSA (Jones, 1994).

Regardless of the dominant acceleration mechanism, we can approximate the acceleration timescale of a CR nucleus with atomic number Z and energy E by a magnetic field of strength B as

$$t_{\rm acc}(E) = \frac{E}{\eta\, B\, Z\, e\, c^2} = \frac{R_L(E)}{\eta\, c}, \tag{2.1}$$

where $R_L(E) = E/ZeBc$ is the CR Larmor radius and $\eta < 1$ is a factor that accounts for the acceleration efficiency. For first-order Fermi acceleration in strong relativistic shocks, this factor is higher the faster the shock and the larger the correlation length of the magnetic field (Gallant &

Achterberg, 1999). In mildly relativistic shocks, η may approach unity (Murase, 2007); however, in ultra-relativistic shocks the efficient return of the particles to the shock front for re-acceleration becomes increasingly difficult, decreasing the acceleration efficiency (*e.g.* Araudo et al., 2015, 2016).

As CRs get accelerated to higher energies, their magnetic rigidity increases, yielding a higher acceleration timescale. When this timescale reaches the light-crossing time of the acceleration region, R/c, the particle cannot be contained in the source and escapes; this dictates the maximum energy to which the source can accelerate an isotope with a certain charge:

$$E_{\rm max} = 10^{20}\,\eta\,Z\left(\frac{B}{1\ {\rm G}}\right)\left(\frac{R}{0.1\ {\rm pc}}\right)\ {\rm eV}. \tag{2.2}$$

This is the minimal condition for the size and magnetic field strength of the sources of UHE protons, and is known as the Hillas criterion (Hillas, 1984).

A nucleus with higher charge $q = Ze$ has a lower magnetic rigidity (defined as E/Z), which makes the acceleration process more efficient, *cf.* Eq. (2.1). This means that if protons can be accelerated to $E^p_{\rm max}$ in a certain magnetic field and source size, a nucleus of atomic number Z will be accelerated in the same source up to energy $E_{\rm max}(Z) = ZE^p_{\rm max}$ due to its higher electric charge. This trend is known as Peters cycle (Peters, 1961) and describes the maximum energy of different nuclei in sources that are optically thin to hadronic interactions. On the contrary, in sources with dense target photon fields, high-energy nuclei will be disrupted before escaping, and the Peters cycle will not apply to the ejected CRs (*cf.* Chapter 3).

2.1.3 Cosmic ray escape and propagation

The escape mechanism of CRs from their acceleration sites depends on the source properties and varies greatly among source classes. In general, estimating the escape timescale of CRs is done by means of numerical models, taking into account the particular source properties such as the geometry of the acceleration region and the magnetic field configuration (*e.g.* Bell et al., 2013; Chen et al., 2016). An analytical approach, however simplistic, can contribute to the generic understanding of the effect of escape on the accelerated CR spectrum. Two extreme scenarios are a purely advective escape on the one hand, and a purely diffusive escape on the other. These will be adopted in Chapter 3 as simplified escape models in blazars, while in this section the discussion will be kept more generic.

In an advection-dominated escape scenario, the CRs are transported by magneto-hydrodynamic (MHD) winds, which represent a bulk motion of the plasma; in the extreme case, these winds are relativistic and the CR escape timescale may therefore be approximated as the light-crossing time of the emission region, $t_{\rm esc} = R/c$. Evidence for such relativistic plasma winds

has been observed in AGNs (Fabian, 2012). This escape mechanism is energy-independent, as we assume that the bulk transport of the plasma is more efficient in transporting the accelerated particles than their own diffusive motion.

A somewhat opposite scenario would be that of a purely diffusive mechanism with no advective transport. Then, the particles in the plasma will be contained longer in the source by the magnetic fields, and their escape is therefore bound by diffusion in the plasma. However, particles close enough to the boundaries of the emitting region can still escape, and this fraction of escaping particles is higher for more energetic particles, since they have larger Larmor radius. A simple modeling of this escape is to assume that the escape rate scales directly with the Larmor radius of the particle (*e.g.* Baerwald et al. (2013); Globus et al. (2015); Zhang et al. (2017)),

$$\frac{c\,t_{esc}}{R} = \frac{R}{R_L(E)} = Z\,\frac{R}{0.01\text{ pc}}\,\frac{B}{1\text{ G}}\left(\frac{E}{10\text{ EeV}}\right)^{-1}, \tag{2.3}$$

assuming $R_L(E) < R$, as implied by Eqs. 2.1 and 2.2. Importantly, since this escape mechanism favors the escape of more energetic CRs, it leads to an emission spectrum that is harder than that produced by the acceleration mechanism. Namely, if inside the source CRs are accelerated to an $E^{-\alpha}$ spectrum, the escaping CR spectrum will scale as $E^{-\alpha+1}$, due to the energy dependence of Eq. (2.3).

As mentioned in Section 2.1.1, the origin of the cutoff of the UHECR spectrum at 100 EeV (Figure 2.1) is not yet clear. On the one hand, Auger fits suggest that this may be the maximum energy of CRs escaping the sources (Aab et al., 2017b). On the other hand, if CRs of higher energies are emitted, the so-called Greisen-Zatsepin-Kuzmin (GZK) effect (Greisen, 1966; Zatsepin & Kuzmin, 1966) can also create such a cutoff in the observed spectrum, depending on the distance of the CR sources. This effect is a suppression of the UHECR flux due to photo-hadronic interactions with the isotropic background photons during propagation in the intergalactic medium. Above a certain threshold, CRs emitted by an extragalactic source will interact with the CMB photons by photo-meson production (*e.g.* $p\gamma \to p\pi$), thus losing energy to secondary mesons (see Section 2.4 for an account of hadronic and electromagnetic CR cooling processes). For protons, that threshold is around 50 EeV, which means that protons detected above this energy must have originated in the local Universe (at 100 EeV, the energy loss length is of the order of 100 Mpc). For heavier nuclei, the threshold for photo-meson is higher and their energy loss is less efficient; however, UHE nuclei are disrupted through photo-disintegration (*cf.* Chapter 3) on the CMB and the extragalactic background light (EBL, Allard, 2012), and are therefore limited to even shorter cosmological distances. For example, a 50 EeV iron-56 nucleus has a mean free path of less than 100 kpc due to photo-disintegration on the CMB.

At energies at the ankle and below, protons lose energy during propagation through electron-

positron pair production on the CMB ($p\gamma \rightarrow p\,e^+e^-$). Compared to photo-meson production, pair production leads to lower energy loss per interaction, yielding a longer mean free path of about 1 Gpc for 10 EeV protons. However, protons below ~ 50 EeV are strongly deflected by the Galactic magnetic fields (Farrar & Sutherland, 2017), therefore losing the directional information necessary to trace back their sources. Moreover, for distant sources, the intergalactic magnetic field (IGMF) may contribute even more strongly to the deflection of UHECRs. Past studies have suggested that 100 EeV protons can be deflected up to 20° by the IGMF on their path from the source to the detector (Sigl et al., 2004), making the detection of extragalactic proton sources difficult, while other models have yielded more optimistic results (Dolag et al., 2004, estimate deflections of less than 1° at the same energy).

2.2 Astrophysical neutrino observations

Several of the difficulties in tracing back the sources of CRs can in theory be averted through observations of cosmic neutrinos. Neutrinos are abundantly produced by high-energy CRs in environments with high radiation density, chiefly as sub-products of photo-pion production (*cf.* Section 2.4). In astrophysical objects with high matter density, neutrinos can also be efficiently produced through nucleus-nucleus interactions between high-energy CRs and the nuclei of the dense medium. Unlike CRs, once produced, neutrinos can escape the source unscathed due to the very low interaction cross sections, and can then travel cosmological distances undeflected by magnetic fields. This means that a detected neutrino points back directly at the position in the sky of its production site. Moreover, the low cross sections of neutrino interactions with matter imply that neutrinos can probe probe as far as the early Universe (the only factor that leads to the cooling of neutrinos on the scale of cosmic travel distances is the adiabatic expansion of the Universe); on the contrary, high-energy photons and CRs are limited to the local Universe (*cf.* Sections 2.1.3 and 2.4). These two advantages can, in theory, open the door to extragalactic neutrino astronomy.

Several neutrino telescopes exist nowadays, and throughout this work we will be basing ourselves in results from IceCube, a km^3 scale underground detector of high-energy neutrinos, located at the geographic South Pole. The reason for such a large detection volume is the extremely low cross section of interactions between neutrinos and matter; in the case of IceCube, the interaction target is the Antarctic ice itself. The detection of a high-energy neutrino in IceCube occurs when the neutrino, while crossing the polar ice, interacts with a nucleus of the ice, producing a charged lepton, most commonly an electron or a muon. This interaction is exceedingly rare, since it occurs via weak interactions. The out-going lepton will carry the momentum of the original neutrino, producing Cherenkov radiation as it crosses the ice. The

detection volume is instrumented with over 5000 digital optical modules (DOMs) intended to detect this emission, which is why the purity of the ice is crucial. Other neutrino telescopes, like the current ANTARES and the future KM3NeT, use the water of the Mediterranean Sea as the target for neutrino interactions.

The signals observed so far in IceCube can be categorized into two "topologies" according to the geometry of the distribution of excited DOMs, and these are *showers* and *muon tracks*. A shower occurs when the neutrino interaction produces an electron or a tau lepton, which then interacts with the electrons in the ice, initiating a spherically-shaped electromagnetic cascade whose photons can be detected by only a few hundred of DOMs in the vicinity of the original interaction. In the cases when this interaction occurs within the fiducial volume of the detector (a case named *starting event*), the deposited energy can be estimated with an uncertainty of only 15% above 10 TeV (Aartsen et al., 2014), through which the energy of the original neutrino can be estimated, under certain assumptions. However, showers provide a weak angular resolution of the neutrino direction of around 10 degrees at most, making it difficult to determine the direction of the source. The other type of event is a track, which has the specificity that it can only be produced by muon neutrinos (ν_μ) via charged current (CC) interactions[1]. The muon produced can travel many kilometers, sometimes crossing the entire detector without being absorbed, ionizing the atoms along its trajectory and producing light that is collected by the DOMs in the vicinity of the track. A track allows for a good determination of the direction of the incident neutrino, down to a resolution of 1 degree. However, the determination of the neutrino energy is more uncertain in these cases, because only a small amount of the neutrino energy is deposited in the detector.

The left panel of Figure 2.2 shows the neutrino spectrum observed by IceCube over the course of eight years, considering only the sub-sample consisting of high-energy starting events above 60 TeV. In blue is represented the estimated flux of *conventional* atmospheric neutrinos, which are produced in the decay of charged pions and kaons from hadronic air showers initiated by CRs. The sub-dominant component shown in green is that from prompt atmospheric neutrinos, also expected to be produced in air showers in the fast decay of charmed mesons. Given that the atmospheric components fall steeply as a function of energy, the observed high-energy flux above 300 TeV cannot be explained by atmospheric neutrinos or muons, and must therefore be of astrophysical origin (Aartsen et al., 2013).

As we can see, the diffuse astrophysical neutrino flux reaches as much as 3×10^{-8} GeV cm^{-2} s^{-1} at 100 TeV in muon neutrinos, or 10^{-7} GeV cm^{-2} s^{-1} in all flavors. If these neutrinos are pro-

[1]Exchange of a W boson with the target nucleus. On the contrary, showers can be produced by electron or tau neutrinos via CC interactions, and by any neutrino flavor via neutral current (NC) interactions, i.e. through the exchange of a Z boson.)

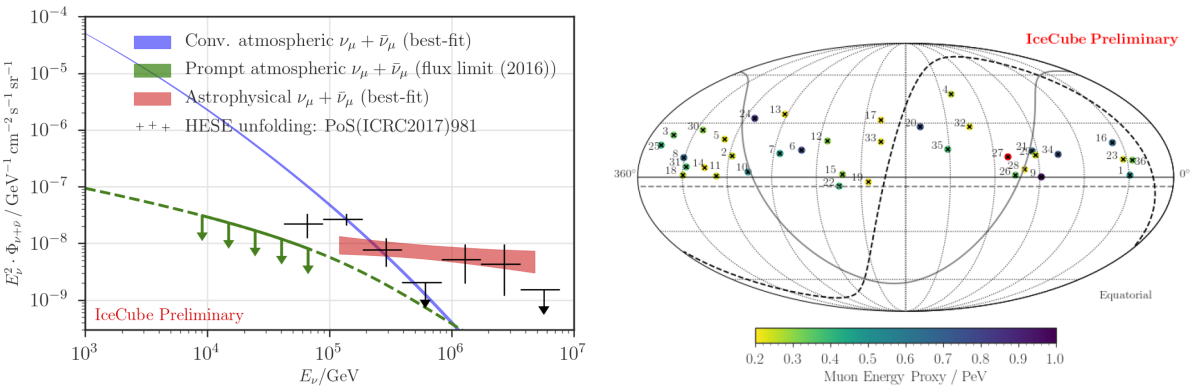

Figure 2.2: *Left:* Energy spectrum of muon neutrinos measured by IceCube in the first 8 years of observations. The black data points represent the sample of high-energy starting events (HESE), while the red band is based only on throughgoing muon tracks. The atmospheric backgrounds are shown as the blue and green curves. *Right:* Sky map of muon neutrinos with energy above 200 TeV observed by IceCube in the same period. Figures taken from Haack & Wiebusch (2018).

duced in lepto-hadronic processes in high-energy sources, then a comparable amount of gamma rays should also be produced in these sources. If those gamma-rays are all released from the source and eventually reach us after cascading down to GeV-TeV energies during propagation, in certain cases this leads to tension with the observed isotropic diffuse gamma-ray background (IGRB) (Murase et al., 2016). As we will see in Section 2.3.1, blazars where neutrinos are efficiently produced should indeed be optically thick to their own gamma-ray emission due to photon annihilation, which can help alleviate this tension.

The flux resulting from such an analysis using throughgoing muon tracks is shown as a red band in the left panel of Figure 2.2. Firstly, by selecting the subset of muon tracks that originate below the horizon, the Earth may be used as a filter to the background of atmospheric muons, which carries a large uncertainty from hadronic shower modeling; moreover, this analysis was performed above 200 TeV, where the contribution from atmospheric neutrinos is also negligible. The red band is obtained under the assumption that the astrophysical flux follows a simple power law in this range, and the spectral index obtained is 2.19 ± 0.10.

Because of the good angular resolution of muon tracks, this event topology is suitable for searching for point sources of astrophysical neutrinos. The neutrino sky after eight years of IceCube observations is shown in the right panel of Figure 2.2, where only throughgoing muons were considered. While theoretically neutrinos should be emitted alongside with photons in photo-hadronic interactions, this neutrino sky map does not correlate with the positions of gamma-ray sources (Palladino & Vissani, 2017; Haack & Wiebusch, 2018). By comparing the

neutrino arrival directions with catalogs of different high-energy source classes, limits have been placed on their contribution to the diffuse flux. For instance, the contribution of observed GRBs to the neutrino flux has been limited to only a few percent (Aartsen et al., 2017b). By the same token, the collection of all observed gamma-ray blazars can contribute only up to 50% of the neutrino flux; any higher contribution is contradicted by the lack of correlations between the neutrino sky and gamma-ray blazar catalogs (Aartsen et al., 2017a).

In the context of this analysis, it is important to note the fact discussed in Section 2.3.2 that most blazars are not resolved as point sources due to their low luminosity or high redshift. While the stacking analysis constrains the contribution of resolved sources (above the orange line in Figure 2.6), the large number of low-luminosity BL Lacs currently undetected may still be responsible for the observed diffuse neutrino flux. In Chapter 4, I will develop the argument that unresolved BL Lacs can indeed explain the high-energy IceCube flux without violating the constraints put by lack of correlations (Palladino et al., 2019), an analysis to which I contributed with an estimation of the neutrino efficiencies of blazars of different luminosities, using the hadronic model described in Chapter 3.

2.3 Introduction to blazars

2.3.1 Multi-wavelength observations and the AGN picture

An AGN is a type of galaxy that possesses an active supermassive black hole in its center, with a mass of millions to billions of times that of the sun. About 10% of these active galaxies emit a relativistic jet of plasma from the galactic core region, depicted in Figure 2.3. The jet has typical bulk Lorentz factors in the order of $\Gamma \gtrsim 10$ and is launched perpendicularly to the plane of the accretion disk. In some AGNs, the jet can maintain its plasma stability over a surprisingly long distance of 10^5 pc, extending well into the intergalactic medium. In the unified AGN model (Urry & Padovani, 1995), blazars are AGNs whose jet is aligned with the direction of observation, which intensifies the observed emission through relativistic effects, often outshining the stellar emission from the entire host galaxy.

Currently, the largest sample of gamma-ray blazars is that observed by the Large Area Telescope (LAT) aboard the Fermi Gamma-ray Satellite, which for over 10 years has continually scanned the entire high-energy sky in a range from 20 MeV to over 300 GeV. In addition, several ground-based gamma-ray telescopes also contribute to the discovery of blazars, such as H.E.S.S (Hinton, 2004), VERITAS (Weekes et al., 2002), MAGIC (Baixeras et al., 2004), and HAWC (Sinnis et al., 2004). Similarly to UHECR telescopes, ground-based gamma-ray telescopes measure the electromagnetic air showers created by gamma rays as they hit the top of

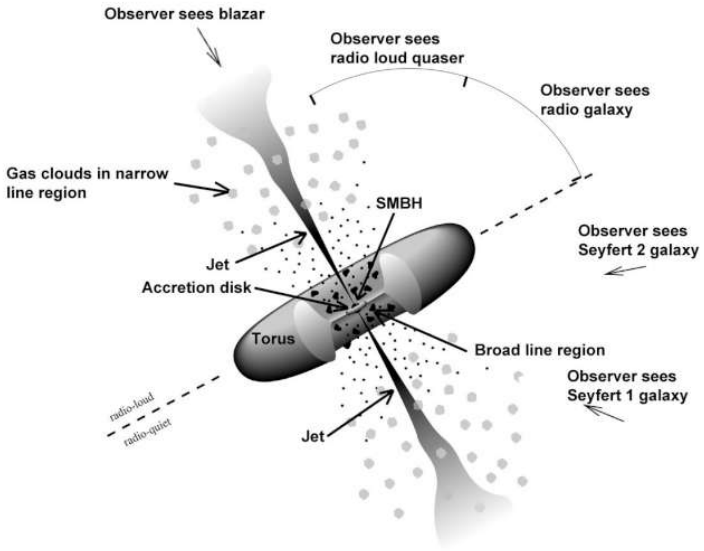

Figure 2.3: Schematic representation of a radio-loud AGN, with indication of the different astronomical classes observed from different viewing angles. Figure taken from Urry & Padovani (1995).

the atmosphere in order to detect Cherenkov radiation produced by the charged particles in these showers. While ground-based telescopes have a more limited field of view compared to space-based observatories, they can measure gamma-rays up to VHEs (reaching 100 TeV in the case of HAWK), a factor of a thousand above the LAT range. As we will see in the following sections, multi-wavelength observations of blazars can constrain crucial source properties, and are used to model the particle interactions taking place inside the jet.

The spectral energy distribution (SED) of blazars is characterized by two broad bands spanning the electromagnetic spectrum from radio to gamma-rays, in some cases up to TeV energies. An example of a simple SED is shown in the left plot of Figure 2.4, representing the SED of Markarian (Mrk) 421, one of the blazars closest to Earth. The low-energy band ranges from radio up to infrared (IR), optical, or even X-rays. The high-energy emission band ranges from X-rays to gamma rays.

Owing to the wide variety of peak frequencies of the blazar emission bands, a sub-classification exists that distinguishes low-, intermediate- and high-synchrotron peaked (LSP, ISP, and HSP) blazars, based on whether their low-energy band peaks below 10^{14} Hz, between 10^{14} and 10^{15} Hz, or above 10^{15} Hz, respectively (Abdo et al., 2010a). From the SED of Mrk 421 we can see that this object is a HSP blazar.

The other elements of AGNs represented in Figure 2.3 are in come cases relevant for the

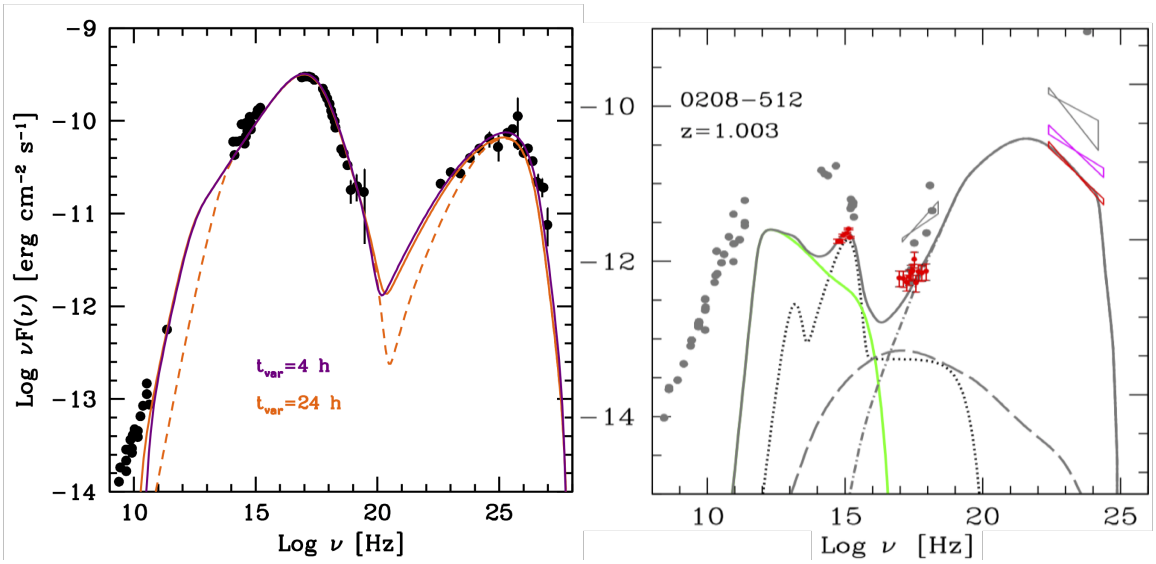

Figure 2.4: *Left:* Multi-wavelength data from the BL Lac Mrk 421 (Abdo et al., 2011) and fitting synchrotron self-Compton (SSC) model curves for different parameter sets obtained by Tavecchio & Ghisellini (2016) (from where the plot is taken). *Right:* Multi-wavelength data from the FSRQ PKS 0208-512 (Abdo, 2010; Ghisellini et al., 2011) and a fitting external Compton (EC) model by Ghisellini et al. (2011) (from where the plot is taken). The points and bands are observations, and the curves represent the model result: the solid curve is the total, and the contributing processes are marked separately: electron synchrotron (green), SSC (long dashed, sub-dominant in this case), EC (dot-dashed). The dotted curve represents the external fields from the torus, disk and corona. Note that the disk contribution is observable as a "big blue bump" in the SED.

processes taking place in the jet. The accretion disk consists of hot plasma spiraling into the black hole. The disk spectrum consists of a complex thermal emission in the optical-UV range, reflecting a multi-temperature profile, while the corona (*i.e.* the outermost layer surrounding the disk) may in some cases be an efficient X-ray emitter. The average total luminosity of the disk has been empirically related to the black hole mass (Ghisellini, 2010), and in some powerful blazars the disk emission is actually observable, as it sticks out from the jet spectrum in the corresponding frequency band. This can be seen, for instance, in the spectrum of the blazar PKS 0208-512, shown in the right panel in Figure 2.4, where a "big blue bump" appears at 10^{15} Hz due to the disk emission, whose theoretical model is shown as the dotted curve. A component from the corona emission is also included in the disk emission model as a soft spectrum extending up to hard X-rays.

Another structure of AGNs is the dust torus (see Figure 2.3), a ring of dust surrounding the

black hole at a typical radius of several light years. In AGNs with bright disks, the disk emission heats up the exposed dust in the torus, which radiates in the infrared. This is represented in Figure 2.4 by the peak at 10^{13} Hz, which in this case has a sub-dominant contribution to the total SED, unlike the disk emission.

Finally, the components marked in Figure 2.3 as *narrow* and *broad line region* (BLR) represent molecular clouds of gas that surround the black hole and the accretion disk. When shined upon by the disk, the gas in the BLR emits spectral lines, the brightest of which is the Lyman-alpha line at 10.2 eV, emitted by hydrogen gas. On the other hand, the gas lying further away in the narrow line region will emit narrow spectral lines. Broad line emission is an important feature that is used to divide blazars into two classes: BL Lacs (a class named after the galaxy BL Lacertae) have absent, or at most faint, broad line emission. That is the case of Mrk 421 (left panel of Figure 2.4). The BLR of BL Lacs is dim or non-existing, and their spectrum is dominated by the non-thermal emission of the jet at all wavelengths. On the other hand, blazars with bright broad lines (above 5 Å equivalent width, Marchã et al., 1996) are called flat-spectrum radio quasars (FSRQs), as is the case of PKS 0208-512. These are high-luminosity blazars with a highly massive black hole, a bright accretion disk, and a large BLR that reprocesses the disk emission. This is why in FSRQs the disk emission is often observable in the overall spectrum, as is some times the IR emission from the dust torus. Most gamma-ray-bright blazars are of the FSRQ type, while most low-luminosity blazars are BL Lacs, as discussed in Section 2.3.2.

When an AGN is observed at large angles from the direction of the jet, the non-thermal emission from the jet is relativistically beamed away from Earth, leading to the observation of a non-blazar galaxy, with a spectrum that extends up to only a few hundreds of keV. In radio-loud galaxies, the jet dominates the radio emission, and its structure can in some cases be resolved by radio telescopes. On the other hand, in radio-quiet AGNs there is no bright jet, which allows for the probing of the host galaxy with radio telescopes, and in some cases it is actually possible to resolve the structure of the core region of the AGN itself (Panessa et al., 2019). As represented in Figure 2.3, in the case of an observation angle that is oblique to the galaxy plane, the BLR of the galaxy core may be visible, leading to the observation of a quasar, the most common type of radio-loud AGN, characterized by both broad and narrow lines. On the contrary, if the observation angle is along the plane of the galaxy, the dust torus may obscure the broad lines emitted in the BLR, leading to the observation of the AGN as a narrow-line AGN. This correspondence between blazars and different types of radio galaxies has been argued by Urry & Padovani (1995), and is known as the AGN unification scheme.

Recalling the blazar classification into HSPs, ISPs and LSPs regarding their synchrotron peak frequency, the observed ISP and HSP blazars are almost all intermediate- and high-frequency peaked BL Lacs (IBLs and HBLs, respectively) and have relatively low gamma-ray luminosities

(Abdo et al., 2010a). On the other hand, the observed FSRQs, most of which are very bright in gamma rays, are typically LSPs (although there are also low-frequency peaked BL Lacs, LBLs). This has led to an empirical relationship between the gamma-ray luminosity of blazars and the peak frequency of their synchrotron emission, a concept known as *blazar sequence* (Fossati et al., 1998). Originally developed based on 33 blazars detected by the EGRET telescope (the predecessor of the Fermi LAT), it has recently been updated (Ghisellini et al., 2017) to include all the blazars with known redshifts from the *Fermi* 3LAC catalog (Ackermann et al., 2015a), currently the most comprehensive catalog of gamma-ray blazars. This trend is represented in Figure 2.5.

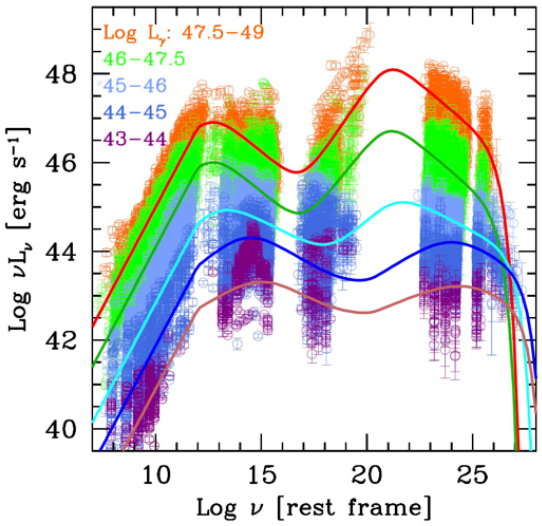

Figure 2.5: Representation of the blazar sequence, taken from Ghisellini (2016). The curves represent the evolution of the average SED along the blazar sequence. The data points represent the observed fluxes of the blazars considered in the sample, at different wavelengths. The color indicates the range of the gamma-ray luminosity bin, whose logarithm is indicated.

The concept of blazar sequence is controversial, and has been contested mainly due to observational biases against high-redshift BL Lacs (*e.g.* Giommi et al., 1999). This bias is due to the poor sensitivity to gamma-ray sources of low luminosity and high redshifts, as discussed in the next section. It provides nonetheless a useful relationship based on the currently known gamma-ray blazars, and it will be used in Chapters 3 and 4 to obtain average SEDs of blazars in different luminosity ranges.

2.3.2 Cosmological blazar evolution

The redshift of a blazar can be identified through spectroscopic analysis, when broad or narrow emission lines are observable in the spectrum; while this is possible for FSRQs, many BL Lacs have featureless SEDs, making a redshift identification not always possible (although indirect methods also exist). Out of the 1500 blazars in the *Fermi* 3LAC catalog, 747 have an identified redshift, which allows the study of their cosmic evolution. In Figure 2.6 are two depictions of a blazar evolution model by Ajello et al. (2012, 2014), which gives the number of sources per unit cosmological volume as a function of the redshift (*redshift distribution*) or gamma-ray luminosity (*luminosity function*). The authors describe the luminosity function with a set of 12 parameters, which are different for BL Lacs and FSRQs. On the left panel we show the result of a Monte Carlo simulation using the luminosity function provided by this model (Palladino et al., 2019). The orange area corresponds to combinations of redshift and luminosity that yield a gamma-ray flux below $\sim 10^{-14}$ erg cm^{-2} s^{-1}, which means those sources are not resolved by *Fermi*. The actual number of points shown (indicated in the figure separately for resolved and unresolved BL Lacs and FSRQs) has been obtained by normalizing the total number of resolved sources (*i.e.* above the orange line) to the total number of blazars in the *Fermi* catalog. We see that most blazars predicted from a theoretical standpoint are in fact low-luminosity, unresolved sources. In particular, only one-seventh of the total number of BL Lacs predicted by evolution models are nowadays resolved as point sources, while for FSRQs that fraction is about one-half. This fact will be relevant in Chapter 4 when we discuss neutrino emission from the entire blazar population. Note that while the unresolved blazars outnumber the resolved ones, the gamma-ray flux is actually dominated by resolved sources, due to the fact that those are on average brighter emitters.

Although most low-luminosity BL Lacs are not resolved as point sources, their gamma-ray emission contributes to the IGRB Ackermann et al. (2015b). The IGRB has been measured by *Fermi* in the range from 100 MeV to 0.8 TeV (Ackermann et al., 2015b), and is obtained by subtracting the contribution of resolved point sources to the total extragalactic gamma-ray background (EGRB). In fact, based on the luminosity function of AGNs, Di Mauro et al. (2014) have concluded that this source class alone can account for the entire IGRB.

By integrating the distribution over redshift, we obtain the blazar density only as a function of luminosity, which is shown in the right panel of Figure 2.6 for BL Lacs and FSRQs separately. We see that the FSRQ population has very high luminosities, tightly concentrated around 10^{48} erg s^{-1}, while BL Lacs are more widely distributed around lower luminosities. On the other hand, the BL Lac distribution is double-peaked in this representation. An analysis of the redshift distribution shows that BL Lacs around the first luminosity peak, dimmer than

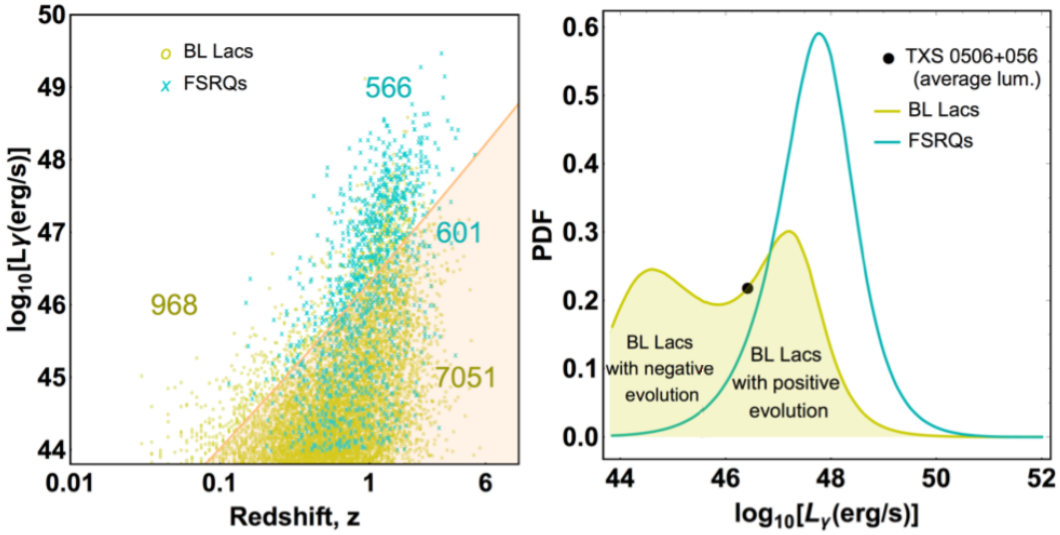

Figure 2.6: *Left:* cosmological evolution of BL Lacs (yellow points) and FSRQs (blue points), according to the luminosity function by Ajello et al. (2012, 2014). The shaded area marks the blazars with flux below the Fermi LAT sensitivity, whose number far surpasses that of resolved blazars. *Right:* redshift-integrated density distribution of BL Lacs (yellow) and FSRQs (blue). BL Lacs are split into two populations (separated by $L_{\rm evo} = 3.5 \times 10^{45}$ erg s^{-1}, indicated in red), while the FSRQ distribution peaks at high luminosities. The black dot represents the average luminosity of the blazar TXS 0506+056 (see Section 2.5 and Chapter 5). Figure adapted from Palladino et al. (2019).

$\sim 10^{46}$ erg s^{-1}, follow a trend of negative cosmological evolution (*i.e.* their density decreases with redshift), while brighter BL Lacs have positive evolution. The existence of different trends in the evolution of blazars of different luminosities may indicate the existence of two distinct populations, which may therefore have different physical properties.

2.4 Radiative interactions in blazar jets

The non-thermal emission of blazars gives insight into the particle interactions taking place in the jet. The first important fact to note is that the relativistic speed of the jet introduces relativistic effects that change the energies and timescales measured by an observer in the rest frame of the shock where the cosmic ray acceleration takes place (henceforth referred to as the *shock rest frame* and indicated by primed quantities) and an observer co-moving with the host galaxy, (which is also the rest frame of the black hole that powers the AGN). Additionally, between the black hole rest frame and the reference frame of an observer at Earth, there are

effects introduced by the redshift z of the host galaxy. These effects may be summarized in three aspects:

- The energy of the emitted radiation appears blue-shifted by the Doppler factor δ, given by
$$\delta(\theta_{\rm obs}) = \frac{1}{\Gamma_{\rm b}(1-\beta\cos\theta_{\rm obs})}, \tag{2.4}$$
where β is the bulk jet speed of the jet in speed-of-light units, $\Gamma_{\rm b}$ is the bulk Lorentz factor of the jet and $\theta_{\rm obs}$ is the angle of observation (a head-on observation would correspond to $\theta_{\rm obs} = 0$, for which $\delta = 2\Gamma_{\rm b}$);

- Relativistic time dilation takes place, which combined with light propagation effects increases the apparent rate of physical processes in the observer frame by a factor of δ relative to the plasma rest frame;

- Due to relativistic beaming, the solid angle Ω of the emission appears contracted by a factor δ^2, which leads to an increase by the same amount in the observed intensity.

- The cosmological drift of the host galaxy leads to the redshift of the emission, which decreases the photon energy by a factor $1+z$ and increases the observed timescales by the same amount.

Note that while the bulk Lorentz factor of the jet $\Gamma_{\rm b}$ is a property of the source, the Doppler factor δ depends on the viewing angle: by Eq. (2.4), the Doppler boost is strongest for head-on observations and decreases as the observation becomes more off-axis. The opening angle of the jet, where most of the emitted radiation is concentrated due to beaming, is well approximated by $\theta = 1/\Gamma_{\rm b} \lesssim 5^\circ$, for which $\delta = \Gamma_{\rm b}$. With this in mind, the effects listed before imply that the observed energy of the emitted radiation is given by $E = E'\,\Gamma_{\rm b}/(1+z)$, and the observed variability timescale by $t_{\rm var} = t'_{\rm var}(1+z)/\Gamma_{\rm b}$. The isotropic-equivalent luminosity of the source inferred from a flux measurement is therefore given by $L = L'\,\Gamma_{\rm b}^4/(1+z)^2$, where a factor of $\Gamma_{\rm b}^2$ is included to account for the beaming effect.

In contrast, the *physical luminosity* is given in terms of the observed luminosity by
$$L_{\rm phys} = L\left(\frac{1+z}{\Gamma_{\rm b}}\right)^2, \tag{2.5}$$
and corresponds to the jet power observed from the black hole frame (without the $\Gamma_{\rm b}^2$ boost that originates from the contraction of the emission solid angle in the rest frame of an observer on Earth, as discussed earlier). The physical luminosity is relevant because it is the rate at which the jet effectively draws power from the black hole, as measured in the black hole's own rest

frame. A potential limit to this power is called the *Eddington luminosity* and corresponds to the maximum power that can be steadily emitted by an accretion flow (in this case, the accretion disk of the black hole) without compromising its structural integrity. Evidently, this stability is maintained by the gravitational pull by the black hole itself, which counteracts the outwards pull of the emitted radiation. This is why the Eddington luminosity depends linearly on the mass of the black hole, $M_{\rm BH}$:

$$L_{\rm Edd} = \frac{G\, M_{\rm BH}\, m_p\, c}{\sigma_T} \approx 10^{46} \left(\frac{M_{\rm BH}}{10^8\, M_\odot}\right) \text{ erg s}^{-1}. \tag{2.6}$$

This argument holds only for steady flows, which means that the Eddington limit can in principle be temporarily exceeded during blazar flares (Sadowski & Narayan, 2015); moreover, the Blanford-Znajek mechanism Blandford & Znajek (1977), which provides an additional means of extracting kinetic energy from the black hole, can also relax the Eddington limit. Notwithstanding these particular considerations, the Eddington luminosity provides a rough limit to the physical luminosity carried away by the AGN jet: $L_{\rm phys} \lesssim L_{\rm Edd}$. This is highly relevant for the radiation modeling of blazars, since it limits the power that is available in electrons and CRs to inject into the radiation zone, L_e and $L_{\rm CR}$, respectively. Of particular interest in the context of lepto-hadronic models (*cf.* Section 2.4.2) is that the Eddington luminosity can put a limit on the *baryonic loading* (or CR loading, $\xi_{\rm CR}$) of the source, given in general by the ratio between the luminosity of injected CRs (protons or nuclei) and electrons:

$$\xi_{\rm CR} = \frac{L_{\rm CR}}{L_e} \sim \frac{L_{\rm CR}}{L_\gamma}, \tag{2.7}$$

where the last scaling holds in the case where the radiation field in the source is in energy equipartition with the electrons ($u'_\gamma \sim u'_e$ where u is the energy density, *cf.* next section).

2.4.1 Leptonic interactions

In most blazar models, the low-energy peak of blazar SEDs originates in synchrotron emission form a population of non-thermal electrons. The energy of the synchrotron peak then depends on the distribution of non-thermal electrons and the magnetic field strength B' in the source. In the monochromatic approximation, where all electrons in the source have the same energy and an isotropic momentum distribution, the total luminosity of the synchrotron emission is given by (see *e.g.* Tavecchio & Ghisellini, 2016):

$$L_\gamma^{\rm syn}(\gamma'_e) = \frac{4}{3}\sigma_T\, c\, u'_B\, \gamma'^2_e\, N_e\, \Gamma_{\rm b}^4, \tag{2.8}$$

where σ_T is the Thomson scattering cross section, N_e is the total number of synchrotron-emitting electrons, $u'_B \sim B'^2$ is the magnetic field energy density in the emission region, $\gamma'_e = E'_e/(m_e c^2)$ is the electron's Lorentz factor in the shock rest frame and $\Gamma_{\rm b}^4$ accounts for the relativistic boost of the emission (for simplicity, the redshift of the source will be neglected throughout this discussion).

The frequency of the synchrotron peak also depends on the magnetic field strength and on the energy of the electrons accelerated in the source. Each electron in the jet radiates a synchrotron spectrum that peaks at a characteristic energy (Blumenthal & Gould, 1970), given in the observer's frame by

$$E_\gamma^{\rm syn} \approx \frac{e\,\hbar\,\Gamma_{\rm b}\,B'\,\gamma_e'^2}{m_e} \sim 5\,\frac{\Gamma_{\rm b}}{10}\,\frac{B'}{10\ {\rm G}}\left(\frac{E'_e}{\rm GeV}\right)^2\ {\rm eV}. \tag{2.9}$$

For a realistic non-thermal electron population, the spectral shape of the synchrotron emission depends on the spectral index of the electron distribution. In the simple case where where the electron distribution follows an unbroken power law of $E'^{-\alpha}$ (between some $\gamma_e'^{\rm min}$ and $\gamma_e'^{\rm max}$), the synchrotron emission spectrum to the left of the peak will scale as $E'^{-(\alpha-1)/2}$ (Longair, 1994). However, in reality synchrotron emission leads to the cooling of the accelerated electrons. If acceleration takes place continuously, synchrotron cooling introduces a break in the electron spectrum above the energy $E_e'^{\rm break}$ for which synchrotron cooling becomes more efficient than acceleration: $t'_{\rm syn}(E_e'^{\rm break}) = t'_{\rm acc}(E_e'^{\rm break})$, where $t'_{\rm acc}$ is the acceleration timescale of electrons given by Eq. (2.1) and $t'_{\rm syn}(E') \sim E'$ is the synchrotron cooling timescale, which can be obtained from Eq. (2.8). Above $E'^{\rm break}$, the electron spectrum will be softer than the one produced by acceleration. As discussed in Section 2.1.2, the maximum energy $E'^{\rm max}$ is in principle obtained by equating $t'_{\rm acc}$ with the light-crossing time of the acceleration zone (by the Hillas criterion) or the cooling timescale if it is lower than the light-crossing time. In the models discussed in this work, we will generally assume that charged particles in the source are indeed continuously accelerated while radiating. If, on the contrary, the radiative processes take place without continuous injection of freshly accelerated particles, the cooling will lead to a cooling cutoff rather than a break.

Depending on the compactness of the emission region in the jet and on the magnetic field strength, the *synchrotron self-absorption* of the electrons may leave a signature in the low-energy emission in the form of a break in the radio spectrum. Below this break, the emission spectrum scales always as $E'^{5/2}$ regardless of the electron distribution, as long as the magnetic field in the source is isotropic (Longair, 1994) and the unabsorbed synchrotron spectrum follows a power law to low energies. More importantly, the energy of this break depends only on the magnetic field strength and the compactness of the emission region, which may be used to constrain these

properties.

The origin of the high-energy peak is in general a more complex topic, as there is a variety of high-energy processes that can be responsible for its emission. In *leptonic* models, gamma rays are the result of the inverse Compton (IC) scattering of low-energy photons by the accelerated electrons. In simple models of BL Lacs, the only available target photons for IC scattering are those from synchrotron emission of the same electron population. This scenario corresponds to a synchrotron self-Compton (SSC) model (Konigl, 1981; Maraschi et al., 1992; Bloom & Marscher, 1996). An example of this scenario is the model depicted in the left plot of Figure 2.4. The ratio between the luminosities of the inverse Compton peak, $L_{\rm C}$ ad the synchrotron peak, $L_{\rm C}$ is called the *Compton dominance*, and in SSC models gives an estimate of the ratio between the electron and magnetic field energy densities (Tavecchio & Ghisellini, 2016):

$$\frac{L_{\rm IC}}{L_{\rm syn}} = \frac{\xi U'^{\rm syn}_{\gamma}}{U'_{\rm B}}, \tag{2.10}$$

where $\xi \approx 1$ for the Thomson regime, when the target photons are soft in the electron rest frame ($\varepsilon_{\rm syn} \ll m_e c^2$), and $\xi < 1$ for the Klein-Nishina (KN) regime, which takes place for hard photons in the electron rest frame ($\varepsilon_{\rm syn} > m_e c^2$). This is because in the Thomson regime the IC cross section is well approximated by σ_T, while in the KN regime the scattering is inelastic and the cross section decreases with the target photon energy.

In some blazars configurations, external fields can be Lorentz-boosted into the jet (see Section 2.4.3) and provide the main target for IC scattering. This model, called external Compton (EC), is often applied to FSRQs, where the external target photon field is provided by the BLR emission (Dermer & Schlickeiser, 1993; Sikora et al., 1994; Ghisellini & Madau, 1996), or by the IR emission from the dust torus (Sikora et al., 1994; Wagner & Witzel, 1995; Sikora et al., 2002). In some BL Lac models, synchrotron emission produced in other regions of the jet may also act as target photons for EC emission (Ghisellini et al., 2005; Georganopoulos & Kazanas, 2003). In any of these cases, because the IC peak depends on the external fields, the synchrotron and IC emission from the same region in the jet are not correlated, which increases the flexibility of the model compared to the SSC scenario, at the cost of a higher number of degrees of freedom. An example of an EC model is shown in the right panel of Figure 2.4, applied to the blazar PKS 0208-512. Here, the main target photons are provided by optical emission from the disk, reprocessed in the BLR, which are up-scattered by electrons in the jet to up to a few GeV, explaining the observed gamma-ray fluxes.

2.4.2 Lepto-hadronic emission

The second distinct class of models is that of *lepto-hadronic* models. In this paradigm, gamma-ray emission originates in interactions of high-energy hadrons with soft target photons in the environment (Mannheim, 1993). In spite of the evidence that heavier nuclei are accelerated to UHEs, most current lepto-hadronic models consider only proton interactions. In this introduction we will therefore limit ourselves to proton interactions, and in Chapter 3 we will consider the effect of heavier nuclei in CR and neutrino emission from blazars, one of the main research questions addressed in this thesis.

At energies slightly above threshold, $p\gamma$ interactions happen through the excitation of Δ resonance:

$$\begin{aligned} p\,\gamma \to \Delta^+ \to & n\,\pi^+ \ (\mathrm{BR} = 2/3) \\ & n \to p\,e^-\,\bar{\nu}_e \\ & \pi^+ \to \mu^+\,\nu_\mu \\ & \qquad \mu^+ \to e^+\,\bar{\nu}_\mu\,\nu_e \\ p\,\gamma \to \Delta^+ \to & p\,\pi^0 \ (\mathrm{BR} = 1/3) \\ & \pi^0 \to \gamma\gamma. \end{aligned} \tag{2.11}$$

The branching ratios (BRs) indicate that 2/3 of the times a charged pion is produced, and 1/3 of the times a neutral pion is produced. In either case, the inelasticity of the interaction is of about $f_\pi = 20\%$, which is the fraction of the proton energy taken by the pion. If a charged pion is produced, four neutrinos are emitted in the subsequent decays, each neutrino taking approximately 5% of the original CR energy. On the other hand, the neutral pion decays into two gamma rays, each with 10% of the original CR energy. The energy ratio of the emitted neutrinos and gamma rays is therefore given by $E_\gamma \approx 2\,E_\nu$, while their average luminosity is given by $L_\gamma \sim L_\pi$. The fact that photo-hadronic interactions produce energetic neutrinos alongside gamma rays makes them interesting if one wishes to explain simultaneously the emission of photons, UHECRs and neutrinos, which cannot be achieved with purely leptonic models.

The center-of-momentum energy necessary to excite the Δ resonance is about $\varepsilon = 150$ MeV. For example, protons with 100 TeV with in the shock rest frame may interact with UV target photons of $E'_\gamma \sim 2$ keV, producing 10 TeV gamma rays and 5 TeV neutrinos. This is one of the process illustrated in Figure 2.7, which shows the photon and neutrino emission predicted by a lepto-hadronic model of the blazar TXS 0506+056 (Rodrigues et al., 2018b, *cf.* Section 2.5 and Chapter 5). Once produced, the neutrinos simply escape the source, and are observed with

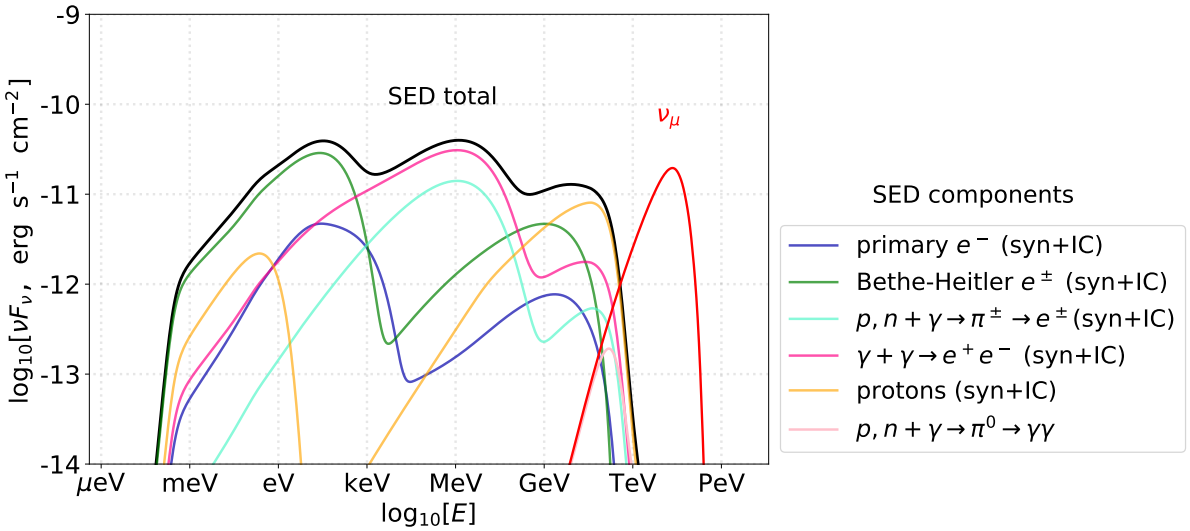

Figure 2.7: Multi-wavelength and neutrino fluxes obtained with a lepto-hadronic model of TXS 0506+056 (Rodrigues et al., 2018b). The different components of the emitted SED are plotted separately. As we can see, in this model the emission from radio up to MeV is dominated by synchrotron by secondary electron pairs, produced by proton Bethe-Heitler pair production ($p\gamma \to pe^+e^-$) and cascades from hadronic interactions, while the gamma ray emission is dominated by proton inverse Compton. Figure courtesy of Shan Gao.

energy $E_\nu \approx 100$ TeV $= \Gamma_\mathrm{b} E'_\nu/(1+z)$. The neutrino flux represented in Figure 2.7 corresponds to the observed flux of muon neutrinos exclusively, assuming a one-third fraction of this flavor after propagation.

The gamma rays that are co-produced through the decay of neutral pions may re-interact through photon annihilation with softer photons in the source, and produce electron-positron pairs:

$$\gamma_\mathrm{VHE}\ \gamma_\mathrm{IR-Xray} \to e^+e^-, \tag{2.12}$$

with a threshold energy of $2m_e \approx 1$ MeV (*e.g.* TeV gamma rays will annihilate with optical photons with energy 1 MeV^2/1 TeV = 1 eV). In turn, the electrons and positrons will cool in the source through synchrotron or IC scattering, creating an electromagnetic cascade that will continue to develop while these processes are efficient, which depends on the source properties. The energy originally released by $p\gamma$ interactions in VHE gamma rays is thus re-distributed to sub-threshold energies, which in the example of Figure 2.7 can be observed as a bump at ∼MeV energies. In FSRQs, where there is an additional soft component from external emission, this process becomes more pronounced and an absorption dip or even a cut-off may appear at GeV

energies (as discussed in depth Section 2.4.4). Therefore, a high neutrino efficiency implies in general a high optical thickness of the source to gamma rays (Dermer et al., 2007; Waxman & Bahcall, 1999), especially given that the cross section for photon annihilation is on average a factor of 10^3 higher than that of $p\gamma$. This means that in general, neutrino-efficient blazar models can be well constrained X-ray observations, while the relationship $L_\gamma \sim L_\nu$ expected from a naive assessment of Eq. (2.11) does not necessarily hold (as we will see for example in the case of TXS 0506+056 in Chapter 5).

Even if a blazar is optically thin to its own gamma emission, photons that escape the blazar into the intergalactic medium can also efficiently annihilate with the soft photons of the EBL, which consists mainly of emission from stars and dust. After the CMB, this is the most energetically dense background in the intergalactic medium, particularly in the range from far infrared (FIR) to UV. During propagation, gamma rays with energy above ∼TeV will annihilate with EBL photons in this range and produce electromagnetic cascades, thus dumping the VHE blazar emission at sub-threshold energies. This effect is more pronounced for high-redshift sources, leading to a cut-off at sub-TeV energies (like in the SED of Figure 2.7, where the gamma rays from photo-meson production are strongly attenuated). However, the cascaded emission will be gradually de-beamed, producing a halo of gamma-ray emission surrounding the point-like source, or becoming more isotropized and thus contributing to the extragalactic gamma-ray background (*cf.* Section 2.3.2).

Besides interacting hadronically through photo-meson production, high-energy protons also interact electromagnetically with the environment radiation through photo-production of electron-positron pairs ($p\gamma \rightarrow pe^+e^-$), also called Bethe-Heitler pair production. The threshold for this process is lower than that of photo-meson production (center-of-mass energy of $2m_e \approx 1$ MeV). Given also that it has a higher cross section and likely more abundant target photons in the astrophysical context, this process can become the dominant cooling process for protons below the photo-meson production threshold. This process is highly relevant in the case depicted in Figure 2.7, where the synchrotron emission from Bethe-Heitler electron-positron pairs far outshines that from primary (*i.e.* injected) electrons (dark blue curve). In fact, in this example the fluxes at all wavelengths are dominated either directly or indirectly by proton-triggered processes

Finally, a class of hadronic models that has not yet been mentioned is that of proton synchrotron (Mucke & Protheroe, 2001). In this case, the high-energy flux form blazars originates in proton synchrotron emission, rather than interactions with target radiation. Compared to electron synchrotron, the energy of the radiation emitted by a proton with the same energy is lower by a factor $(m_e/m_p)^3 \sim 5 \times 10^9$ (*cf.* Eq. 2.9, noting that the Lorentz factor scales with the mass of the particle). This is evident from Figure 2.7, where the proton synchrotron emission peaks at only 0.1 eV, with a magnetic field of 1 G and maximum proton energy of 70 TeV in the

shock rest frame (see Rodrigues et al., 2018b). Is is therefore evident that proton-synchrotron blazar models require stronger magnetic fields in the jet and higher maximum proton energies compared to other lepto-hadronic models in order to explain gamma-ray emission.

2.4.3 One-zone models

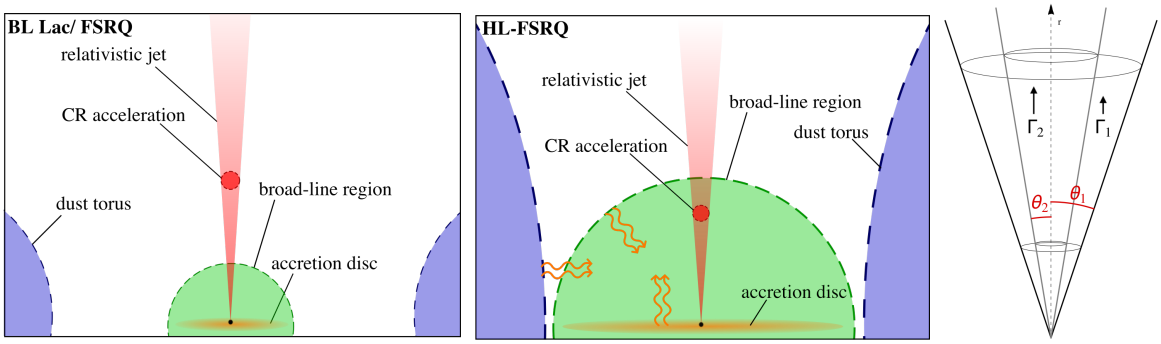

Figure 2.8: *Left and middle:* geometry of a one-zone BL Lac model *(left)* and an external field FSRQ model *(middle)*. In the case of FSRQs, if the acceleration region in the jet (blob, dark red) is lying inside the BLR (green), the bright emission from the external components are relativistically boosted into the blob, where it will contribute to the non-thermal interactions. Figures taken from Rodrigues et al. (2018a). *Right:* representation of the jet in a spine-layer model. The spine is traveling with a bulk Lorentz factor Γ_2 that is higher than that of the surrounding layer, Γ_1. Figure adapted from Sikora et al. (2016).

The model behind the emission spectrum of Figure 2.7 is a one-zone model, the simplest in terms of geometry. In this model, the emitting region in the jet is assumed to be a spherical blob, represented in the left panel of Figure 2.8. This blob is moving along the jet with the bulk Lorentz factor of the plasma, Γ_b. The observed variability timescale of the emission, in some cases as small as minutes, can constrain the size of the blob through the generic argument that photons escape the blob at light speed, and therefore the blob size must roughly correspond to the distance traveled by a photon in the variability timescale, t_var, transformed into the shock rest frame:

$$R'_\mathrm{blob} \lesssim 3 \times 10^{16} \, \frac{t_\mathrm{var}}{1 \text{ day}} \, \frac{\Gamma}{10} \text{ cm}, \tag{2.13}$$

where the Lorentz factor accounts for the relativistic transformation $t'_\mathrm{var} = \Gamma t_\mathrm{var}$, due to the motion of the blob towards the observer. The injection of accelerated particles and the radiation emitted from non-thermal processes are usually considered to fill the entire volume of the blob, $V_\mathrm{blob} = (4\pi/3)\, R'^3_\mathrm{blob}$, and to be isotropic and homogeneous in the shock rest frame. The photon

density in the blob can then be estimated from the observed luminosity:

$$U'_\gamma = \frac{L_\gamma}{4\pi R'^2_{\rm blob} c \Gamma^4} \approx 0.4 \frac{L_\gamma}{10^{48}\ {\rm erg\ s^{-1}}} \left(\frac{t_{\rm var}}{1\ {\rm day}}\right)^{-2} \left(\frac{\Gamma_{\rm blob}}{10}\right)^{-4} {\rm erg\ cm^{-3}}. \tag{2.14}$$

For comparison, 0.4 erg cm^{-3} is the energy density of a 10 G magnetic field.

If the blob lies at a distance $R_{\rm diss}$ from the black hole (the dissipation radius of the jet), the size of the blob in the direction perpendicular to the jet direction is given by $R'_{\rm blob} = R_{\rm diss} \sin\theta_{\rm jet} \approx R_{\rm diss}\theta_{\rm jet}$, considering a small opening angle $\theta_{\rm jet}$ that is constant along the jet. This size does not get contracted to an observer in the black hole frame, because it is measured perpendicularly to the flow. Assuming an opening angle $\theta_{\rm jet} \approx 1/\Gamma_{\rm b}$ (*cf.* beginning of Section 2.4), this yields a dissipation radius of

$$R_{\rm diss} \sim 3 \times 10^{17} \left(\frac{t_{\rm var}}{1\ {\rm day}}\right) \left(\frac{\Gamma_{\rm blob}}{10}\right)^2 {\rm cm}, \tag{2.15}$$

or ~ 0.1 pc. As the blob propagates along the jet at relativistic speeds it expands, which may lead to the adiabatic cooling of the plasma, unless thermal energy is continuously injected into the plasma via certain dissipative processes. In the most dramatic case, the radial expansion of the outer shell occurs at the speed of light; this will introduce an energy-independent cooling process with a rate given by the (inverse) size of the region, $t'^{-1}_{\rm adia} = 10^{-6}\ {\rm s^{-1}} = c/(3 \times 10^{16}\ {\rm cm})$.

One-zone models have been applied extensively in the past both within the leptonic and the lepto-hadronic paradigms to explain the steady-state multi-wavelength emission of a variety of blazars (*e.g.* Boettcher et al., 2013). Lepto-hadronic models in particular have been used to estimate neutrino emission from particular blazars (such as TXS 0506+056, see Section 2.5) as well as the diffuse flux emitted by an blazar population (Murase et al., 2014).

2.4.4 External field models

An aspect of blazar geometry that must be taken into account in the modeling of FSRQs is the size and shape of the BLR and the dust torus, and the position of the emission region in the jet within these external radiation fields. This brings us to the class of external field models, where the photon emission from components of the AGN outside the jet (*cf.* Section 2.3.1) contribute to the non-thermal interactions in the jet. There is a positive correlation between the disk luminosity $L_{\rm disk}$ and the gamma-ray luminosity of the jet L_γ (Maraschi & Tavecchio, 2003), which is one of the reasons external fields are more relevant in gamma-ray-bright blazars.

The BLR is often considered to be a relatively thin layer of clouds at a distance $R_{\rm BLR}$ from the black hole (*e.g.* Tavecchio & Ghisellini, 2008), as represented by the green dashed line in Figure 2.8. On the other hand, the geometry of the dust torus is in general less constrained,

but can be assessed indirectly from optical observations of non-blazar radio galaxies (Chiaberge et al., 1999); evidence suggests that FSRQs should have thick circum-nuclear tori, while in most BL Lacs the torus should be thin or inexistent (blue regions in the middle and left panels of Figure 2.8, respectively.)

Both the radius of the BLR and of the dust torus have been observed to scale with the disk luminosity (Greene & Ho, 2005; Kaspi et al., 2005; Cleary et al., 2007), a trend that is often parametrized in phenomenological studies (Ghisellini & Tavecchio, 2008; Murase et al., 2014) through a simplified relationship:

$$R_{\rm BLR} = 10^{17} \left(\frac{L_{\rm disk}}{10^{45}\ {\rm erg\ s^{-1}}} \right)^{1/2} \ {\rm cm} \tag{2.16}$$

$$R_{\rm DT} = 2.5 \times 10^{18} \left(\frac{L_{\rm disk}}{10^{45}\ {\rm erg\ s^{-1}}} \right)^{1/2} \ {\rm cm} \tag{2.17}$$

(a relationship that can be understood if we assume that the disk luminosity scales with its surface area). Comparing these relationships with Eq. (2.15), we see that in bright FSRQs the emission region of the jet may be lying inside the volume covered by the external fields, $R_{\rm diss} < R_{\rm BLR} < R_{\rm DT}$. Therefore these external photons may propagate into the jet, which motivates EC models for FSRQs, as discussed earlier. This situation is represented in a simplified geometry in the middle panel of Figure 2.8.

The BLR reprocesses the disk radiation and re-emits it as both broad lines and a thermal continuum. These components are represented in the left panel of Figure 2.9 in the black hole frame. The black peaks represent the H I and He II Lyα lines, often the most prominent ones in the spectrum (Greene & Ho, 2005). The total luminosity emitted in broad lines is parametrized as a fraction of the disk luminosity, $L_{\rm BL} = f_{\rm cov}^{\rm BL} L_{\rm disk}$, where the covering factor $f_{\rm cov}^{\rm BL} \approx 10\%$ (Kaspi et al., 2005; Liu & Bai, 2006).

The green curve in Figure 2.9 is a universal fit of the thermal continuum (Elvis et al., 1994), usually the sub-dominant component of the BLR emission. The portion of this emission that is of concern for radiation modeling of the jet is only the fraction which is isotropized through Thomson scattering in the BLR electrons, while the components traveling away from the disk (in the same direction as the jet) will appear redshifted in the jet frame and its contribution to non-thermal interactions will be suppressed. The fraction of the BLR continuum that is scattered back depends on the optical thickness of the BLR to Thomson scattering, $\tau_{\rm T}$, which depends on the electron density and size of the BLR. The value of $\tau_{\rm T}$ may in principle vary from 10^{-4} to 10^{-2} (Blandford & Levinson, 1995). The dust torus, on the other hand, emits a single thermal spectrum in the IR range, represented in blue in Figure 2.9, whose total luminosity is also a fraction of the disk luminosity, $L_{\rm DT} = f_{\rm cov}^{\rm DT} L_{\rm disk}$. The torus covering factor $f_{\rm cov}^{\rm DT}$ of

most blazars is not well known, but in previous modelings a value of 50% is often assumed (*e.g.* Murase et al., 2014).

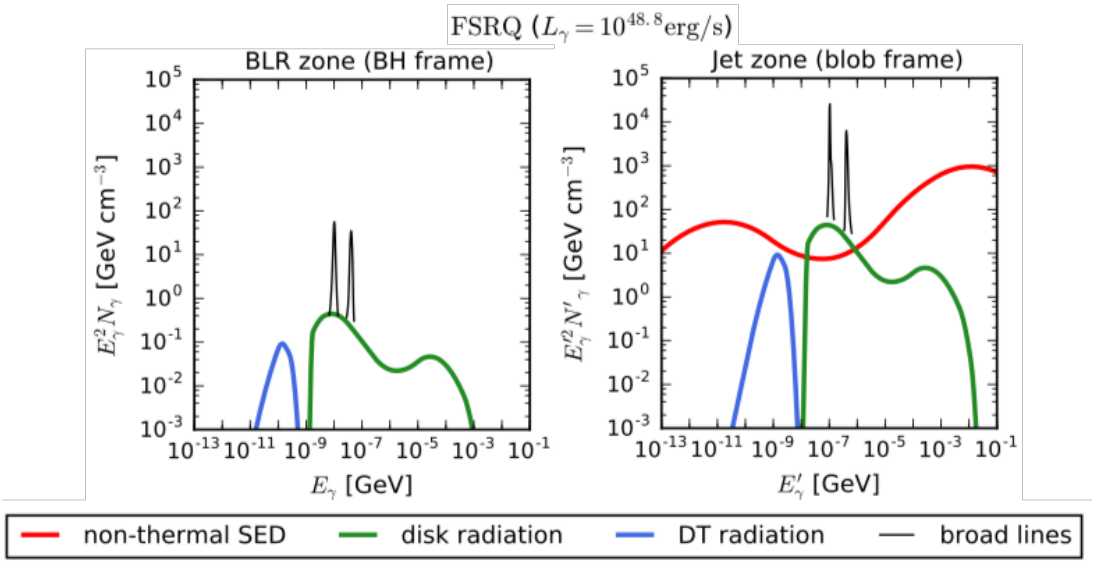

Figure 2.9: Energy spectrum of the external fields of an FSRQ in the black hole frame *(left)* and in the jet frame, together with the non-thermal SED from the jet itself *(right)*. Figure taken from Rodrigues et al. (2018a).

Given the luminosities $L_{\rm BLR}$ and $L_{\rm DT}$, and assuming the emission fills a spherical region with the respective radii, the energy density of these photon fields can be obtained by $U_{\rm BLR,DT} = L_{\rm BLR,DT}/(4\pi R^2_{\rm BLR,DT} c)$, since the photons escape through a surface area of $4\pi R^2_{\rm BLR,DT}$ with an escape speed of c. In the rest frame of the jet blob, this radiation density appears boosted due to Doppler boosting and to the Lorentz contraction of the blob in one direction as it travels within these fields. We therefore obtain:

$$U'_{\rm ext} = \Gamma^2_{\rm blob}\, U_{\rm ext} = \frac{f_{\rm ext}\,\Gamma^2_{\rm blob}}{4\pi R^2_{\rm ext}\, c}\, L_{\rm disk}, \tag{2.18}$$

where the subscript 'ext' refers to an external component; either the BLR continuum, the broad lines and the torus emission, and $f_{\rm ext}$ corresponds to the fraction of the disk emission that is re-processed by this component, respectively $f^{\rm BL}_{\rm cov}$, $\tau_{\rm T}$ and $f^{\rm DT}_{\rm cov}$. The radiation density spectrum boosted into the jet frame is shown in the right plot of Figure 2.9. We can see that the energy of the emission is shifted by a factor $\Gamma_{\rm blob} = 10$, and the energy density of the photon field boosted by a factor $\Gamma^2_{\rm blob}$, with the BLR photons reaching a number density of 10^9 cm^{-3} at 100 eV, dominating over the emission from the non-thermal particles in the jet itself.

In Figure 2.10 we explore the effects of an external field on the photo-hadronic processes in a

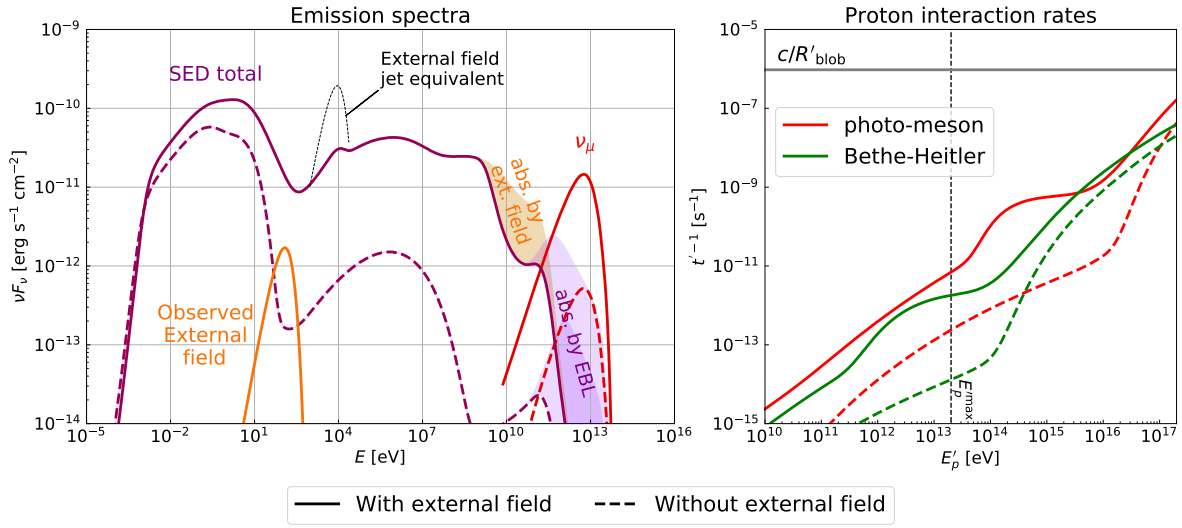

Figure 2.10: Representation of the effect of external fields in a lepto-hadronic blazar model. *Left:* photon and neutrino fluxes with (solid) and without (dashed) the presence of the external thermal field. The orange bump is the thermal field in the observer's frame (out-shined by the jet emission). The black dotted bump is the Lorentz-boosted external field as 'seen' in the shock rest frame (shown here in the observer's frame like the SED). The shaded areas show the effect of gamma-ray absorption due to the same external field during propagation in the broad line region (orange) and on the EBL during propagation in the intergalactic medium (violet). *Right:* Corresponding proton interaction rates in the jet frame of photo-meson production (red) and cooling through Bethe-Heitler pair production (green). Plots obtained with the AM^3 code (Gao et al. (2017), see Chapter 5).

blazar jet. In the left panel we represent the SED and neutrino emission with the inclusion of an external field (solid curves) and without it (dashed curves), while maintaining all the other parameters such as the maximum energy and injection power of protons and electrons, and the magnetic field strength. The spectrum of the external field is represented in orange as seen in the observer's frame. The temperature is 5×10^5 K, which represents well the peak of the continuum emission from the BLR of most FSRQs (Greene & Ho, 2005). A more detailed treatment would have to also take into account the broad lines, which typically dominate the BLR emission, an aspect that will be fully addressed in Chapter 5. We also represent as a black dotted curve the effective flux and energy after boosting into the jet frame[2] according to Eq. (2.18). Note that this component is not observable, since it is directed from the BLR into the jet, in the opposite

[2]The factor $\gamma^2 = 100$ between the frequency of the orange and black bumps in Figure 2.10 is due to the boost by a factor Γ from the black hole frame to the jet frame, and the additional boost by the same factor Γ due to the fact that the spectra are represented in the observer's frame.

direction of the observer (*cf.* Figure 2.3), while the non-thermal SED is considered isotropic in the jet frame.

In the case where no external fields are present, the SED (dashed purple curve) is dominated by the first peak from synchrotron emission, since photo-hadronic interactions of the accelerated protons are inefficient due to the low density of target photons. For the same reason, the emitted neutrino fluxes are low (red dashed curve). On the other hand, as we can see by the solid curves, the inclusion of the external field leads to an increase in X-ray to gamma-ray emission from proton-initiated cascades, and to efficient neutrino production.

The processes at play can be better understood by analyzing the right panel of Figure 2.10, where we plot the energy loss rate of protons through photo-meson production (red) and the proton cooling rate through Bethe-Heitler pair production (green) with and without the presence of the external field. The way to interpret this type of plot is to realize that the higher the inverse timescale for a process, the more efficient it is; therefore, at any given energy, the highest value in the plot corresponds roughly to the dominating process. The optical depth to a process is given by $\tau(E'_p) = t'^{-1}(E'_p)R'_{\rm blob}/c$, which is why $c/R'_{\rm blob}$ is plotted for reference as the gray line. The maximum energy of the protons in this particular example is indicated by the vertical line.

Firstly, we can see the direct effect of the external field serving as a target for both interactions, introducing a bump in the interaction rates at the respective threshold energy (*cf.* the energy of these bumps with the frequency of the black dotted bump in the left panel). The protons in this particular case have sufficient energy to interact directly with the external field through Bethe-Heitler pair production, but not through photo-meson production. The efficient Bethe-Heitler interactions lead to extensive electromagnetic cascades in the source, supported by synchrotron and IC from the secondary electron-position pairs. These cascades lead to an increase in the photon densities across the spectrum, which subsequently increase the efficiency of photo-meson production, as can be seen in the timescale plot: at the maximum proton energy, the photo-meson rate an order of magnitude higher due to this cascade (red solid curve) compared to the case with no external field (red dashed curve).

The increase in the optical thickness to photo-meson production leads to the higher neutrino flux, as shown in the left panel. At the same time, the presence of this field implies that the escaping gamma rays will be further attenuated as they cross the BLR, an effect that has been recently discussed in several works, particularly in the context of the blazar TXS 0506+056 (Murase et al., 2018; Reimer et al., 2018; Rodrigues et al., 2018b). The 100 eV photons of the thermal field annihilate with gamma-rays of $E_\gamma \sim (10^{12}/100)$ eV $= 10$ GeV, leading to the attenuation dip shown as an orange shaded area. Finally, for blazars at high enough redshift, gamma rays harder than 1 TeV (produced directly by photo-meson interactions) are attenuated through annihilation with EBL photons, as indicated by the purple shaded region.

Murase et al. (2014) have included the effect of external fields in the estimation of neutrino production by lepto-hadronic interactions in the jet; their conclusion was that the transport of these external fields into the jet, especially the broad lines, can contributes significantly to neutrino production in high-luminosity FSRQs. The model presented in Chapter 3 attempts to expand this kind of approach by numerically computing for the first time the effect of heavier nuclei (rather than only protons) in CR and neutrino emission from blazars. Furthermore, we take into account explicitly the effect of the propagation of UHECRs through the BLR and dust torus, after escaping the jet through different possible escape mechanisms.

2.4.5 Other configurations

A blazar model with more complex geometry is the spine-layer model (Tavecchio et al., 2014). Here, the non-thermal emission takes place in two separate zones moving with different relativistic speeds along the jet, as represented schematically in the right plot of Figure 2.8. The spine is the internal core of the jet, which moves with a Lorentz factor $\Gamma_s \sim 15 - 20$, and an outer layer moving slower with $\Gamma_s \sim 3 - 5$. Such a structure is supported by evidence from simulations and observations of BL Lac jets (*e.g.* Kovalev et al., 2007). The radiation produced by non-thermal electrons in the layer, with energy density U_l, is 'seen' in the frame of the spine with a boosted density of $U'_l = U_l \Gamma^2_{rel}$, where $\Gamma_{rel} = \Gamma_s \Gamma_l (1 - \beta_s \beta_s)$ is the relative Lorentz factor between the two zones given by relativistic velocity addition. This effects provides a dense target photon field for hadronic interactions and external Compton scattering by electrons and protons accelerated in the spine. The emission form the spine is the dominating contribution to the observed SED, due to the large Lorentz boost in the observer's frame, while the layer emission is sub-dominant in the observer's frame.

If the spine is loaded with protons, neutrino production is more efficient than in one-zone models due to the boost in density of the synchrotron emission (which is the target field for $p\gamma$), which lowers the necessary power injected in protons. Additionally, the neutrinos emitted peak at lower energies than in a one-zone scenario, because the low-energy, sub-threshold synchrotron photons produced in the sheath are blue-shifted when seen in the spine rest frame, becoming eligible to interact with protons which in a one-zone scenario would have too low energies to interact. This provides the model with additional flexibility to explain observations (see *e.g.* the case of TXS 0506+056 in Section 2.5); however, this comes at the cost of additional parameters compared to the simplest, one-zone model. That is also the case of the external-field scenario discussed before; in fact, the two models are similar in that the target photon field is provided by a process external to the main radiation zone. The main difference is the nature of the external radiation, which in this case is non-thermal.

Another type of two-zone model is the compact core model (Gao et al., 2018). In this case, inside the blob there exists a smaller core with higher radiation density, and both regions travel along the jet with the same speed. The low-frequency peak originates in synchrotron emission from a population of electrons accelerated in the large blob, while the GeV emission originates in the smaller core zone, where the high photon density leads to efficient inverse-Compton scattering by an electron population and, in a lepto-hadronic scenario, the protons co-accelerated in the jet will efficiently produce neutrinos through $p\gamma$ interactions. As discussed in the next section, this model was found to explain the 2017 neutrino event from TXS while describing the simultaneous observational data.

A highly relevant aspect of blazar modeling that can potentially challenge conventional geometries is the variability of the emission in the time domain. The light curves of blazars (temporal variation of their emitted flux) are often complex and exhibit occasional periods of violent flaring that can last from days to several months (Nalewajko, 2013). During these flares, the luminosity of the emission may rise by orders of magnitude, sometimes in a matter of hours, and in some cases the frequency of the emission peaks also shifts considerably (see *e.g.* the observations by Catanese et al. (1997) of a flare of Mrk 501). In between these flares, blazars radiate in a quiescent state, with lower gamma-ray luminosity and less variability. In addition, while in HBLs the gamma-ray and X-ray fluxes are typically highly correlated, in many cases the behavior at different wavelengths does not necessarily correlate; an extreme case is the BL Lac object PKS 2155-304 has been observed to undergo violent gamma-ray flares without significant simultaneous variation of the optical or X-ray fluxes (Aharonian, 2009), a case commonly known as an "orphan gamma-ray flare".

The transition from a quiescent to a flaring state may often times be explained by one-zone models through a temporary increase in the injection luminosity of non-thermal electrons, L'_e, or protons, L'_p, into the radiation zone (*cf.* Section 2.5). However, explaining in detail the behavior of some complex light curves is usually beyond such simplified models; for example, the size of the blob required by Eq. (2.13) to explain a minute-scale variability would be extremely small, which in some cases can hinder the acceleration of non-thermal particles to the energy required to explain the emission (Gao et al., 2018), or lead to photon densities in the emission zone that are too high, turn the emission region opaque to gamma rays (Blandford & Levinson, 1995). Alternatively, if we insist on a large blob size, an extremely high Lorentz factor of the jet is often necessary to explain the observed variability. Spine-sheath models partially alleviate this tension through the additional Doppler boost between layer and spine. Other models have also been developed that include different emission regions traveling along the jet at different speeds (Georganopoulos & Kazanas, 2003) or with different orientations, such as in helical jet models (Villata & Raiteri, 1999). Other classes of models that can explain emission variability

are internal shock (Boettcher & Dermer, 2010) and multi-collision models (Spada et al., 2001), where the radiative processes take place in the shock between plasma shells along the jet, like in a GRB fireball scenario. Another example of a model that can naturally predict flares is the star-in-jet model (Giannios et al., 2009), where a red giant crosses the blazar jet, providing a compact region for proton-proton interactions and leading to a gamma-ray flare.

In spite of the importance of reproducing blazar emission in the time domain, in this work we focus on the understanding of multi-wavelength, neutrino and CR emission from blazars, while the study of time-dependent behavior necessarily implies a different focus.

2.5 The 2017 neutrino flare of blazar TXS 0506+056

In September 2017 a 250 TeV muon neutrino left a track in the IceCube detector that allowed to identify its arrival direction as consistent with the position of a gamma-ray source, blazar TXS 0506+056, referred to henceforth as TXS. Due to the triggering of the IceCube alert system (Aartsen et al., 2017c), the source was followed up by telescopes and was observed to be in a state of flaring activity in the radio, optical and gamma-ray bands, a state that lasted several months Aartsen et al. (2018b). Given the temporal and directional coincidence between the neutrino observation and the electromagnetic flare, this was considered the first multi-messenger observation of neutrinos and photons. In the meanwhile, the possibility of an uncorrelated coincidence (for example if the neutrino had originated in a hidden, far-away source in the same direction as TXS) has been excluded at the level of 3.5 σ Aartsen et al. (2018b) on the basis of the temporal and spatial coincidence.

Upon this discovery, through a search in the archival data, IceCube has found a signal, in 2014-15, of approximately 13 neutrinos in excess to the expected (atmospheric) background, distributed over a six-month period from a direction compatible with the same source. Curiously, this neutrino signal was not accompanied by enhanced radio, optical, or gamma-ray activity (Aartsen et al., 2019). Such an "orphan" neutrino flare is challenging from a phenomenological perspective, since as discussed in the previous section, in a neutrino-efficient source there should be activity in the gamma-ray or X-ray bandwidths simultaneously to the neutrino emission, due to the photo-hadronic processes and the cascades that develop thereafter. In Chapter 5 we will investigate this 'orphan' historic neutrino flare of TXS, and in this section we will introduce the current theoretical understanding of this source based on the 2017 flare.

The 2017 event immediately captured the attention of the blazar community, and several lepto-hadronic models have since been applied to explain the neutrino event and the simultaneous multi-wavelength data. One class of these models is based on $p\gamma$ interactions, as introduced in the previous section (Gao et al., 2018; Cerruti et al., 2018a; Zhang et al., 2018; Keivani et al.,

2018; Ansoldi et al., 2018; Sahakyan, 2018; Gokus et al., 2018). Another class of models were also proposed that explain the neutrino and photon emission from proton interactions with matter (proton-proton models Liu et al., 2018; Sahakyan, 2018), analogous to the star-in-jet scenario described in the previous section.

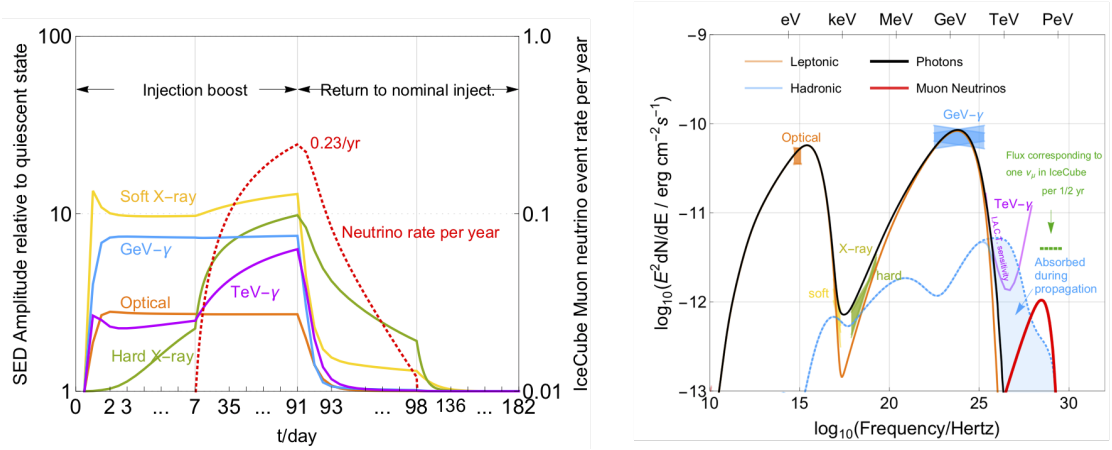

Figure 2.11: Multi-wavelength and neutrino fluxes estimated by the hybrid model of Gao et al. (2018) for the 2017 flare of TXS 0506+056. *Left:* Time dependency of the multi-wave fluxes (solid curves) and neutrinos (dotted curve). *Right:* Spectral energy distributions (orange: leptonic emission; blue: sub-dominant hadronic emission; red: neutrinos). Data points represent observed fluxes during the flare. The dashed horizontal green line is the expected level and energy range of the emitted neutrino flux necessary to produce one muon neutrino event in IceCube over a 180 day period.

The first conclusion that emerges from these studies is that in a one-zone model, the fluxes observed during the flare in the GeV-TeV range cannot be explained by photo-hadronic interactions without overshooting the fluxes of X-rays and TeV gamma rays due to hadronic emission and subsequent cascade (*e.g.* Gao et al., 2018; Keivani et al., 2018; Cerruti et al., 2018b). This immediately presents a limitation to the explanation of the neutrino event, which should require a high enough neutrino flux from the source during the flare.

Given these constraints, the remaining possibility is a kind of hybrid model, where the SED is described purely by leptonic emission, while hadronic interactions of co-accelerated protons can explain some of the expected neutrino emission without the subsequent hadronic cascade exceeding the observed fluxes. In modeling the SED, however, a purely SSC scenario is excluded by the peak frequencies of the observed emission, and therefore other geometries are necessary to explain the emission. Moreover, any purely leptonic model does not account for neutrino

emission.

Gao et al. (2018) have found that a two-zone model can explain the multi-wavelength data, namely the compact core model introduced in the last section. In the left panel of Figure 2.11 we show the predicted emission spectrum of photons and neutrinos together with the multi-wavelength fluxes observed during the flare (Aartsen et al., 2018b). The low-frequency peak originates in synchrotron emission from electrons accelerated in the blob, while the GeV emission originates in the core zone, where the high photon density leads to efficient inverse-Compton scattering by electrons and efficient neutrino production by protons. The predicted neutrino flux corresponds to 0.27 muon neutrino events per year above 120 TeV in IceCube, while any higher neutrino fluxes would lead to brighter emission from the cascading of hadronic gamma rays (blue curve), overshooting the observed levels of X-ray and TeV gamma rays during the flare (yellow, green and violet data points).

In the hybrid model by Keivani et al. (2018), on the other hand, the leptonic emission originates in EC scattering, where an external field, such as from scattered disk emission, serves as the target photon field. Such a scenario may be particularly justified in the case of TXS, since there is now evidence that it may in fact be a masqueraded BL Lac (Padovani et al., 2019). This means that while the broad lines are not observed, leading it to be classified observationally as a BL Lac, it may in fact possess strong external fields in a BLR, corresponding in physical terms to an FSRQ.

A general limitation of these hybrid models in explaining the TXS flare is the high value of L_p/L_e required during the flare. In some cases, the proton luminosity required exceeds the Eddington luminosity of the source (see Section 2.3.1) by several orders of magnitude. A proton synchrotron model should require a lower proton luminosity, which could alleviate the tension with the Eddington limit. As shown by Cerruti et al. (2018b), the multi-wavelength emission can in fact be explained in this paradigm, and the lower proton luminosity required does indeed respect the Eddington limit of TXS. However, the emitted neutrinos would have energies in excess of 1 EeV, which is in tension with the energy of a few hundreds of TeV of the observed muon track.

Another model that has also been applied by Ansoldi et al. (2018) to TXS is the spine-sheath model, introduced in Section 2.4.5. In this case, the authors conclude that the protons may be accelerated to an energy in the range between 100 TeV and 1 EeV in the shock rest frame without violating X-ray constraints. Moreover, due to the efficient neutrino production in the target field provided by the outer layer of the jet, the proton injection luminosity necessary to emit a neutrino flux consistent with the IceCube event is up to only a few times the Eddington limit.

Finally, in terms of the time profile of the flare, both a compact core and an external field

model can explain the transition from the quiescent to the flaring state of the blazar through a time-varying injection of protons and electrons which ramps up and then decays, thus leading to a flare observed across the electromagnetic spectrum, as shown in the right panel of Figure 2.11 for the case of the compact core model of Gao et al. (2018). While the electrons injected in the larger blob would lead to a fast increase in optical, soft X-rays and gamma-rays, the same increase in proton injection in the core leads to a slower build-up of neutrinos and hard X-rays, due to the longer cooling timescale of the hadronic processes.

Although none of the proposed models reproduce a neutrino flux level high enough to explain the observed event, Strotjohann et al. (2018) have shown that this fact cannot on its own exclude these models. In fact, assuming that the 2017 TXS flare is not unique but rather belongs to a distribution of flares from a population of similar sources, the one observed event does not represent the emission levels of TXS alone. This corresponds to the idea of Eddington bias, which is a mis-representation of the characteristics of an entire population of sources based on the observation of a small number of events. Indeed, based on the blazar model introduced in Chapter 3 and on the cosmological evolution of blazars, we will show in Chapter 4 that there should be hundreds of sources capable of producing similar flares to that of TXS.

Drawing from the successes of the 2017 flare modeling, in Chapter 5 we investigate whether the orphan flare of TXS in 2014–15 may also have originated in hadronic interactions in TXS (Rodrigues et al., 2018b). We calculate self-consistently the neutrino and photon emission from lepto-hadronic interactions in three different models: a one-zone model, a compact core model, and an external photon field model, in light of the recent evidence, mentioned previously, that TXS may in fact be an FSRQ (Padovani et al., 2019). My contribution to this project was the development of the external field model. This, simultaneously with Reimer et al. (2018), was the first multi-messenger modeling of the orphan neutrino flare of TXS. We conclude that at most two of the observed 13 neutrinos may have been produced in a photo-hadronic event in TXS without violating the few existing multi-wavelength observations of the same period.

Chapter 3

Beyond protons: blazars with heavier nuclei

In this chapter we present a new model of CR interactions in blazar jets, including isotopes heavier than protons, first published in Rodrigues et al. (2018a). I developed this model based on the numerical code NeuCosmA (Baerwald et al., 2012), which had previously been applied only to GRBs (Hummer et al., 2012; Boncioli et al., 2017; Biehl et al., 2017). This implied adapting the numerical algorithm to make it suitable for the description of blazar flares. I also implemented a three-zone model for the case of FSRQs, in order to properly describe the hadronic processes that occur to UHECRs that escape the jet, during their propagation in the BLR and dust torus radiation fields.

This is the first study to numerically calculate the effect of the blazar photon fields in all relevant photo-nuclear processes including neutrino production. The importance of this work lies in the fact that we now have measurements of the chemical composition of UHECRs that indicate an average composition that is heavier than protons, as discussed in Section 2.1.1. A model of how such nuclei interact in blazars is therefore necessary for making accurate predictions. Moreover, it allows us to accurately quantify how the neutrino emission of blazars depends on the composition of the accelerated CRs, thus providing insight into the connection between these two messengers.

3.1 A photo-hadronic model including nuclei

As discussed in the introduction, protons accelerated in the source will interact through photo-meson production until they have been cooled to energies below the interaction threshold. Additionally, UHECR nuclei heavier than protons can also interact with radiation fields through photo-disintegration. This refers to the splitting of the nuclei into smaller fragments through interaction with the environment photons, most often by splitting off neutrons, protons, and alpha particles: ${}^{A}_{Z}X + \gamma \rightarrow {}^{A'}_{Z'}Y + \cdots$, with $A' < A$ and $Z' \leq Z$.

The isotopes considered in this work are shown in the nuclear chart of Figure 3.1, where the X axis represents the number of neutrons of the isotope and the Y axis represents the number

of protons. In Appendix B we detail the most important aspects of the implementation of the nuclear system, including the criteria for the selection of isotopes. Besides protons and neutrons (lower-left corner of Figure 3.1), we include isotopes with mass number up to $A = 55$, as well as iron-56 (upper-right corner).

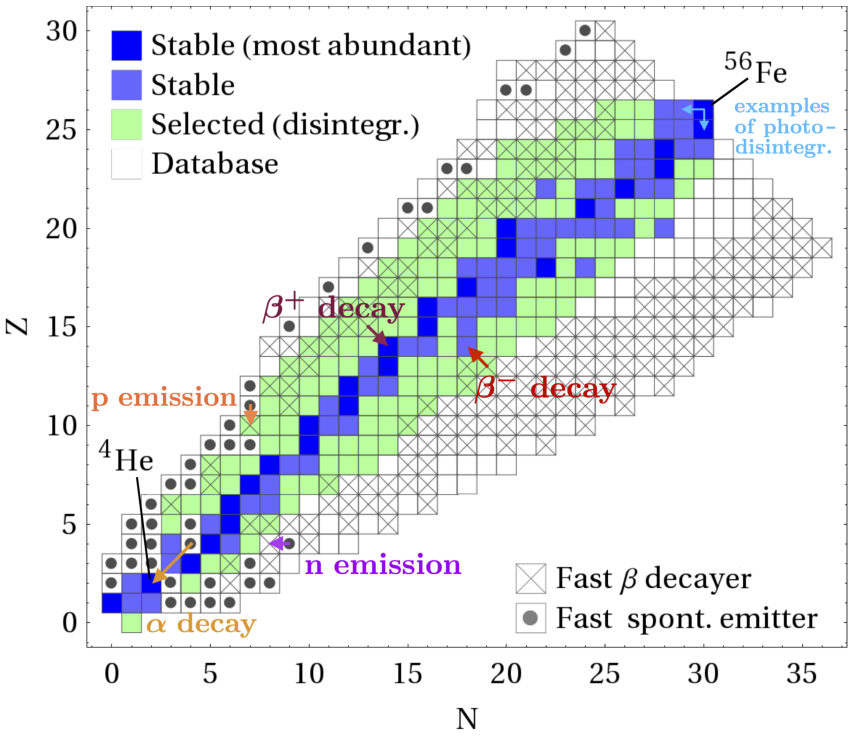

Figure 3.1: Nuclear isotopes considered in this study as a function of neutron number N and proton number Z. Examples of common decay processes are illustrated by colored arrows. Fast spontaneous emitters are marked by dots. As explained in Appendix B, in the simulation these isotopes are replaced automatically with their daughter nuclei, while the colored isotopes are included explicitly in the numerical system. The blue isotopes are stable, while the green ones are unstable $\beta^{\pm}$ emitters. Fast $\beta^{\pm}$ emitters are marked with crosses. The fast decay of these isotopes (those that have been selected) can potentially contribute to the neutrino emission from the jet. Figure adapted from Rodrigues et al. (2018a).

The leading contribution to photo-disintegration is the giant dipole resonance, an excitation generated by a photon above 8 MeV in the nucleus rest frame. Two examples of possible photo-disintegration paths for ion-56 are illustrated in Figure 3.1 as light blue arrows, corresponding to the cases where a proton or a neutron are ejected from a iron-56 nucleus (only the daughter nucleus is indicated, and not the corresponding ejected nucleons). In some cases, the daughter

nuclei of photo-disintegration are unstable, but can nonetheless however have relatively long lifetimes in the shock rest frame due to their high Lorentz factors.

Eventually, unstable elements decay either by spontaneously emitting nucleons or alpha particles, or through beta decay. Beta decay is particularly relevant for neutrino production:

$$ {}^{A}_{Z}X \rightarrow {}^{A}_{Z\pm1}X' + e^{\mp} + \nu, \tag{3.1} $$

where $\nu = \nu_e$ for positron emission and $\nu = \bar{\nu_e}$ for electron emission. On the contrary, spontaneous emission does not produce neutrinos. These spontaneous processes are also represented as colored arrows in Figure 3.1.

The isotopes represented in blue, situated mostly along the diagonal of the diagram, are stable, and therefore can only be transformed into other isotopes through interactions with the photon field. Spontaneous decay processes tend to create isotopes that are closer to this "valley of stability". On the contrary, photo-nuclear interactions create lighter elements that can be unstable, leading to a "widening" of the cascade (in the direction perpendicular to the $Z = N$ diagonal) (Boncioli et al., 2017). For that reason, denser target photons fields for photo-hadronic interactions tend to generate "wider" and more extensive cascades in blazar jets, as will be shown in Section 3.2.

In order to simulate the interactions between the different particle species and the photon and magnetic fields in the jet, we solve the system of coupled partial differential equations (PDEs) that describe the evolution of the population of each particle species. This calculation is done using the NeuCosmA code (Baerwald et al., 2012), which calculates the density spectrum $E'^2 dN'_{\rm i}/dE'$ of all hadronic species i (namely protons, neutrons and heavier nuclei), which is done by numerically integrating the system of the PDEs describing the evolution of each of these spectra. Some details about this method are given in Appendix B. Other works (Boncioli et al., 2017; Biehl et al., 2017) have deployed a similar approach, based on the NeuCosmA code, to the study of hadronic interactions in GRBs.

At a first stage, the PDE system is set up including all the relevant hadronic species. As detailed in Appendix B, the selection of the isotopes is done using a recursive algorithm that follows all possible disintegration and decay paths starting from the heaviest isotope. The isotopes shown in green and blue in Figure 3.1 are those selected to be included explicitly in the PDE system, while the remaining isotopes are produced rarely and do not contribute significantly to the calculation. The isotopes marked with a cross are beta decayers with short enough lifetime that their decay can be relevant for neutrino emission in a timescale comparable to the light-crossing time of the source. Finally, the isotopes marked with a dot are spontaneous emitters with very short lifetimes, and are replaced automatically with their daughters in the PDE

system. The isotopes marked with a cross are beta emitters, and therefore their contribution to the neutrino spectrum is accounted for; on the contrary, those marked with a dot are spontaneous emitters of nucleons or alpha particles, and do not produce neutrinos.

Besides secondary CRs, the hadronic interactions produce mesons (mostly pions, *cf.* Eq. (2.11), but also charmed mesons, like neutral and charged kaons). The decay of these mesons produces neutrinos, whose emission spectrum is also calculated. As mentioned previously, electron neutrinos are also emitted in the beta decay of unstable nuclei, which is taken into account. Because the photon field is taken from observations, the emission of electrons and photons from meson decay is not relevant for the simulation. This means that mesons do not feed back into the PDE system; their density spectra are calculated with the only purpose of estimating the emitted neutrino spectrum.

The interaction rates depend on the magnetic and photon fields in the jet, both considered homogeneous and constant throughout the duration of the simulation. This differs from other photo-hadronic models in that the photon field is not computed self-consistently from the hadronic interactions; instead, the photon density spectrum in the jet, $E'^2 dN'_\gamma/dE'$, is calculated from an SED taken from observations (according to Eq. (2.14)). This static photon field is assumed to be maintained by a population of co-accelerated electrons, which are not included in the calculation. On the one hand, this feature of the model excludes the ability to predict the multi-wavelength emission from blazars; on the other hand, by fixing the photon field from observations, we have direct insight into the efficiency of hadronic interactions in the jet: the geometry of the emission zone is the only additional parameter that affects the interaction timescales. It is important to note that the feedback of hadronic photons to the SED is a non-trivial matter (*cf. e.g.* Figure 2.7), as they may in some cases out-shine the leptonic emission itself (especially in the X-ray range between the two peaks, as in the example shown previously in Figure 2.11). A fully self-consistent hadronic model with heavier nuclei including photon emission is left as a topic for a future work.

For the photo-meson production off nuclei, we assume that the interaction rate scales proportionally to the mass number A of the nucleus. This simplistic treatment was used in other works (Hümmer et al., 2010; Boncioli et al., 2017; Biehl et al., 2017), although more refined treatments may be appropriate for a more accurate description of the neutrino emission from such sources.

The geometry of low-luminosity blazars is based on the one-zone model introduced in Section 2.4.3, where we assume a fixed blob size of $R'_{\rm blob} = 3 \times 10^{16}$ cm, and a bulk Lorentz factor of the blob of $\Gamma = 10$. In Section 3.5, we will consider additionally the effect of external radiation fields characteristic of FSRQs, as discussed in Section 2.4.4.

Rather than a quiescent state where the particle densities in the jet are constant in time, we simulate a hadronic flare of duration $t'_{\rm flare} = R'_{\rm blob}/c = 10^6$ s. A primary isotope is injected in the

blob at a constant rate for the duration of the flare, which means its density in the source grows linearly with time (except at high energies, where under some circumstances the nuclei interact and cool, *cf.* Section 3.2). After $t_{\rm flare}$, the injection comes to a halt and the energy density of the primary CRs fall exponentially to zero due to the adiabatic cooling from the expansion of the blob. During this entire process, the numerical solver is used to integrate the PDE system, and the fluence of CRs and neutrinos that escapes the blob is integrated (*cf.* Appendix B). The particular CR escape mechanism (see Section 2.1.3) will only be discussed in Section 3.4, and until then we will focus on neutrino emission, which is approximately independent from the CR escape rates.

The magnetic field inside the blob is determined by fixing the value of the "magnetic loading" of the source,

$$\epsilon_B = u'_{\rm B}/u'_\gamma \;=\; u'_{\rm B}\left(\frac{L'_\gamma}{4\pi\, R'^2_{\rm blob}\, c}\right)^{-1} \;\sim\; 10^{-2}. \tag{3.2}$$

This yields a magnetic field strength $B' = \sqrt{8\pi u'_B}$ that scales with the gamma-ray luminosity of the source (as $L_\gamma^{1/2}$).

3.2 Cascading of nuclei in the jet

We will now show the effect of the non-thermal photon field on hadronic and nuclear interactions by simulating the injection of iron-56 in a blazar jet using the methods described in the previous section. The first scenario we will discuss is a low-luminosity blazar that is optically thin to photo-nuclear interactions, representative of what we call the “nuclear survival” regime, and the second scenario is a high-luminosity jet that is optically thick to photo-nuclear interactions at the maximum energy, representing the “nuclear cascade” regime. In both examples, the blazar SED is obtained from the (original) blazar sequence (Fossati et al., 1998, see Section 2.3.1), for the respective gamma-ray luminosity bin.

Nuclear survival regime– The example in Figure 3.2 illustrates the injection of a non-thermal spectrum of CRs consisting purely of iron-56, into a jet with observed gamma-ray luminosity $L_\gamma = 10^{44.6}$ erg s^{-1}, defined henceforth in the range $0.1 - 100$ GeV, roughly corresponding to the *Fermi* observation range. The SED considered in the simulation is shown in the upper-left panel, where we represent the photon energy density spectrum in the jet, given in the shock rest frame. The label of the y axis, E'^2N', is shorthand notation for $E^2(dN/dE)$. The corresponding interaction rate plot (*cf.* Figure 2.10 and discussion thereafter) is shown in the upper-right panel. Following the discussion in Section 2.1.2, the maximum energy of the accelerated iron nuclei, $E'_{\rm max} \sim$ EeV, is given by the energy at which the inverse acceleration

timescale (black line) becomes smaller than the light-crossing time of the source (gray line), according to Eq. (2.2), which is not the case for optically thick sources, as discussed below. The acceleration efficiency will be considered to be $\eta = 1$ throughout this chapter.

We can see that the source is optically thin to disintegration at the maximum energy ($\tau_{\rm disint} \lesssim 0.1$) and even more optically thin to other cooling processes like photo-meson production. This is due to the low luminosity of the target photons to hadronic processes. As an example, the energy of the target photons to photo-meson production is marked with an arrow in the SED plot, for nuclei at maximum energy. The low optical thickness of the source to photo-disintegration means that once injected, an iron nucleus will most often survive any interactions; this leads to the isotope density chart in the lower left panel, which shows that the primary species (iron-56) is still the most abundant isotope in the source at the end of the flare. Only a few other isotopes are also present in the source from the few photo-disintegration interactions that take place, but they are several orders of magnitude less abundant. In the lower right panel we can see the energy spectra of the densities in the source of different hadronic species, where it is clear that after the flare the iron spectrum is still unaffected by photo-hadronic interactions, compared to the injection spectrum. As detailed in Appendix B, the shape of the injection spectrum is a power law with an exponential cutoff at the maximum energy. The small offset observed between the injection spectrum and the density of iron-56 after the flare is simply due to adiabatic cooling, assumed to have a constant timescale given by the light-crossing time of the blob (*cf.* Section 3.4).

Nuclear cascade regime– In Figure 3.3 we illustrate in a similar fashion the results of the model for iron-56 injection in a high-luminosity blazar with $L'_\gamma = 10^{48.8}$ erg s^{-1}, whose SED is shown in the upper-left panel. Note that only interactions with the jet radiation are considered, without additional components from the BLR or the dust torus. As can be seen in the upper right plot, the acceleration is more efficient in this case due to the higher magnetic field (*cf.* Section 3.1); however, above 1 EeV the iron nuclei photo-disintegrate more efficiently than they are accelerated. We therefore take that energy as the maximum injection energy, since higher energies would be unrealistic. In general (see Appendix B), in high-energy sources where CR cooling is efficient, we consider the maximum energy to be that for which the sum of the energy loss rates from all cooling processes starts dominating over the acceleration rate. From 100 PeV to 1 EeV, the source is optically thick to photo-disintegration ($\tau_{\rm disint} \lesssim 100$). This results in efficient disintegration of the accelerated particles into secondary nuclei, which creates a cascade of lighter isotopes, as shown in the lower left panel. The most abundant secondary isotopes (purple boxes) are neighbors of the primary nucleus and light fragments like nucleons and α particles. Contrary to other environments like GRBs (Biehl et al., 2017), the disintegration rate strongly increases with energy, which is due to the steep slope of the low-energy peak of the

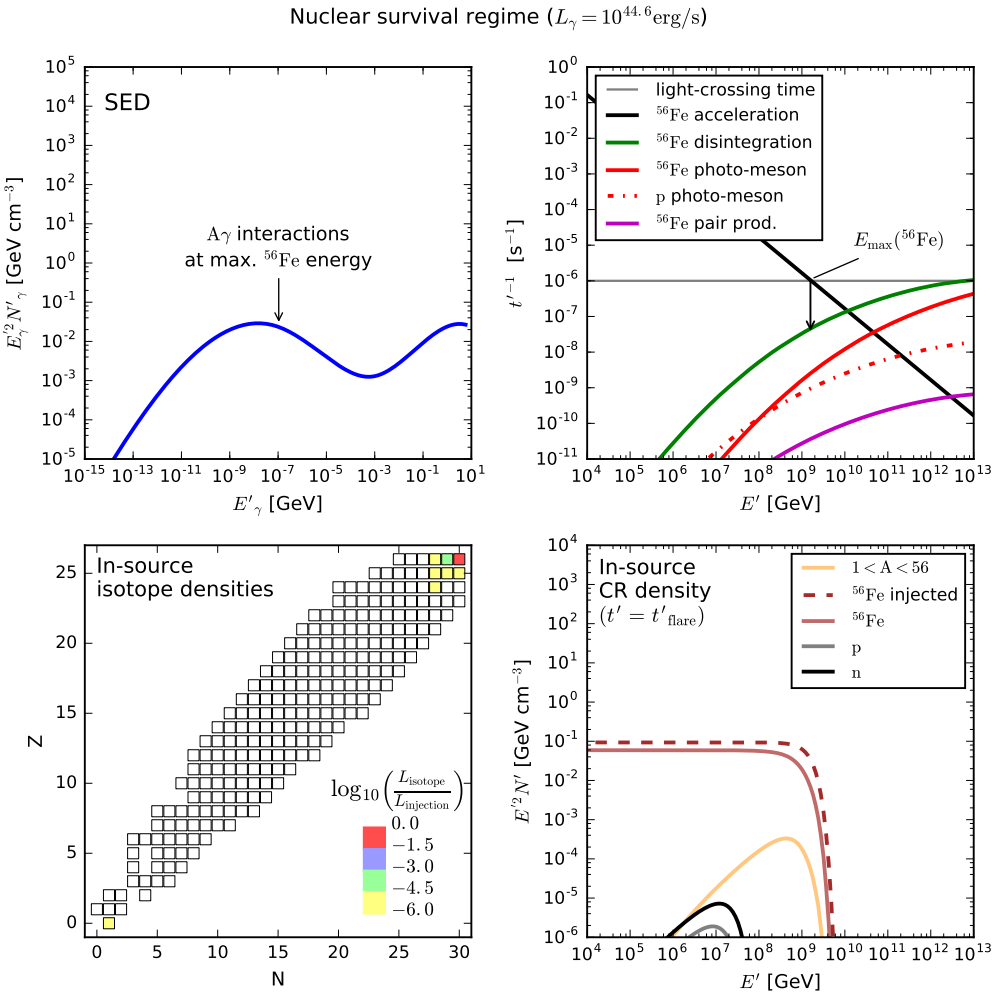

Figure 3.2: Example of the nuclear survival regime. Simulation of the injection of iron-56 in a jet with $L_\gamma = 10^{44.6}$ erg s^{-1}. (*Upper left*) time-independent SED; (*upper right*) time scales (interaction rates) of relevant radiative processes and light-crossing time; (*lower left*) energy densities of nuclear isotopes, normalized to the injection luminosity; (*lower right*) spectra of baryons in the source after light-crossing time (given in the jet rest frame). The optical thickness to $A\gamma$ interactions τ at $E_{\rm max}$ is illustrated using black arrows. Figure taken from Rodrigues et al. (2018a).

SED. As can be seen from the lower right panel, after the flare the CR density is dominated by secondary nuclei from photo-disintegration above 100 PeV, while at lower energies the iron-56 spectrum is unmodified, as expected.

Note how iron-56 photo-disintegration dominates over photo-meson production at all energies, hindering neutrino production directly by iron nuclei (while the secondary CRs may produce neutrinos at lower energies, *cf.* Section 3.3). This fact had in fact already been demonstrated by Anchordoqui et al. (2008). However, this depends on the shape of the SED, and in other sources photo-meson production by nuclei may dominate in some energy ranges.

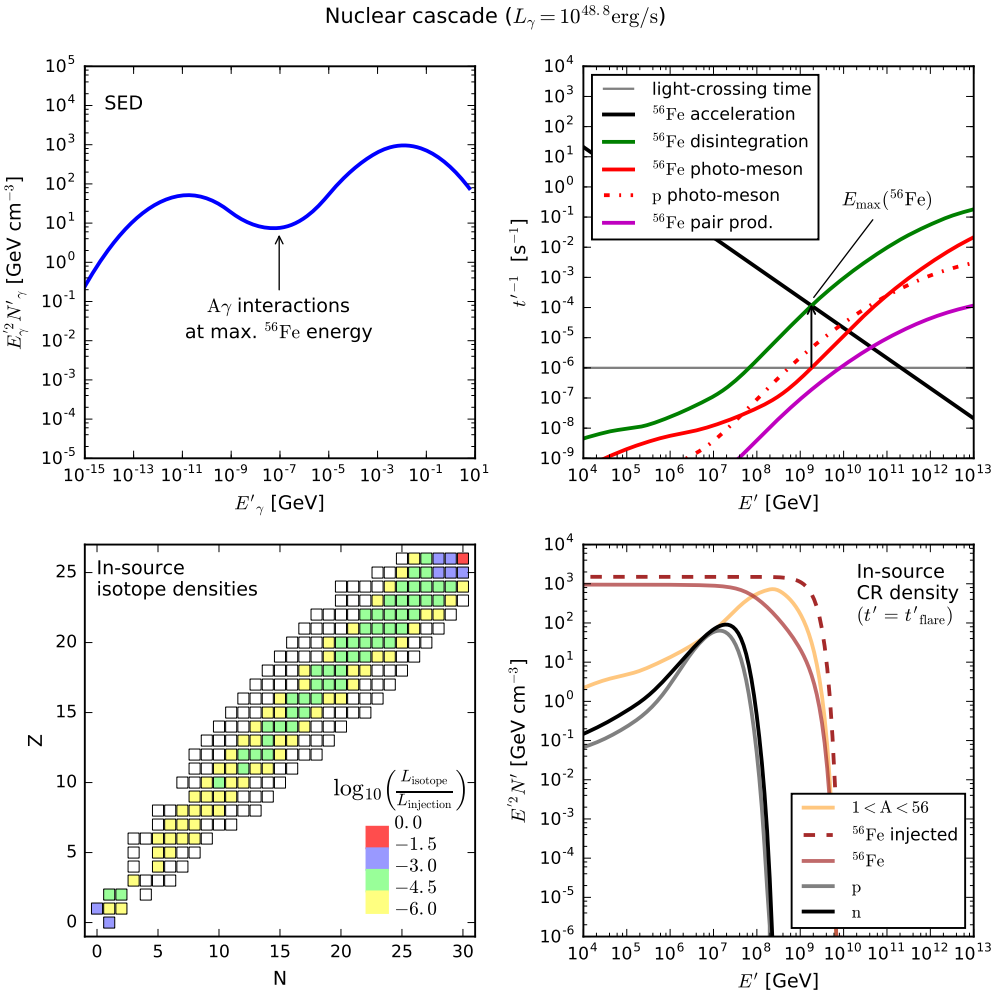

Figure 3.3: Example of the nuclear cascade regime (jet with $L_\gamma = 10^{48.76}$ erg s^{-1}), for a pure iron-56 injection. See caption of Figure 3.2 for details. Figure taken from Rodrigues et al. (2018a).

Parameter space scan– By scanning the parameter space of $R'_{\rm blob}$ and L_γ, we can find the range of parameters where either of the regimes occurs. We define the two regimes based on the optical thickness to photo-nuclear interactions at the highest energy of the injected CRs. The result is shown in Figure 3.4 for three different injection isotopes: protons, helium-4 and iron-56. The shape of the SED is assumed to follow the blazar sequence. From Figure 3.4 we see that by varying L_γ along a fixed line of $R'_{\rm blob}$, we can control the level of disintegration and, as we will see in the next sections, the balance between the efficiency in neutrino and UHECR emission. The same effect can be achieved by varying $R'_{\rm blob}$ for any fixed value of L_γ.

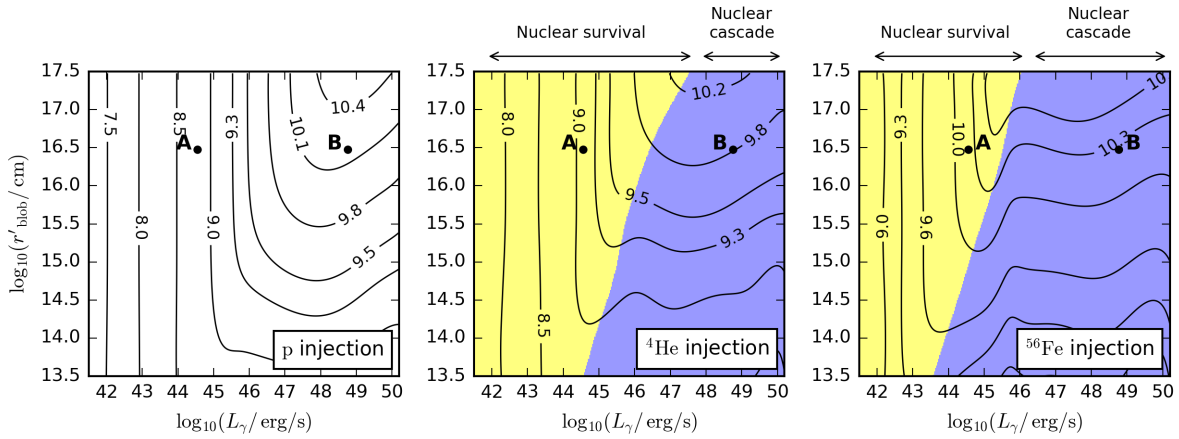

Figure 3.4: The location of the two nuclear regimes in the parameter space of jet luminosity and blob size values, for the injection of three different isotopes. The contours indicate the maximum injection energy in the black hole frame, $\log_{10}(E_{\rm max}/\,{\rm GeV})$. In the center and right-hand panels, the two colors represent the two nuclear regimes, while for protons this does not apply. The points marked A and B represent the examples shown in Figure 3.2 and Figure 3.3, respectively. Figure taken from Rodrigues et al. (2018a).

Also plotted in Figure 3.4 are the contours of the maximum injected CR energy, and we can clearly see that larger blob sizes tend to allow for higher acceleration energies. In general, at low luminosities the maximum energy increases linearly with luminosity due to the corresponding increase in magnetic field, which boosts the acceleration efficiency; a change in blob size does not change the maximum energy simply due to the corresponding decrease in magnetic field, given by Eq. (3.2). On the contrary, at the highest luminosities the maximum energy is limited by hadronic interactions, as explained previously. This introduces a non-trivial shape to the maximum-energy contours. The central and right-hand side panels of Figure 3.4 reveal that the maximum energy is actually achieved in a narrow region in the transition between the two regimes. In this region, acceleration is efficient enough to produce UHECRs, while energy losses

from photo-disintegration are not too strong.

3.3 Neutrino emission

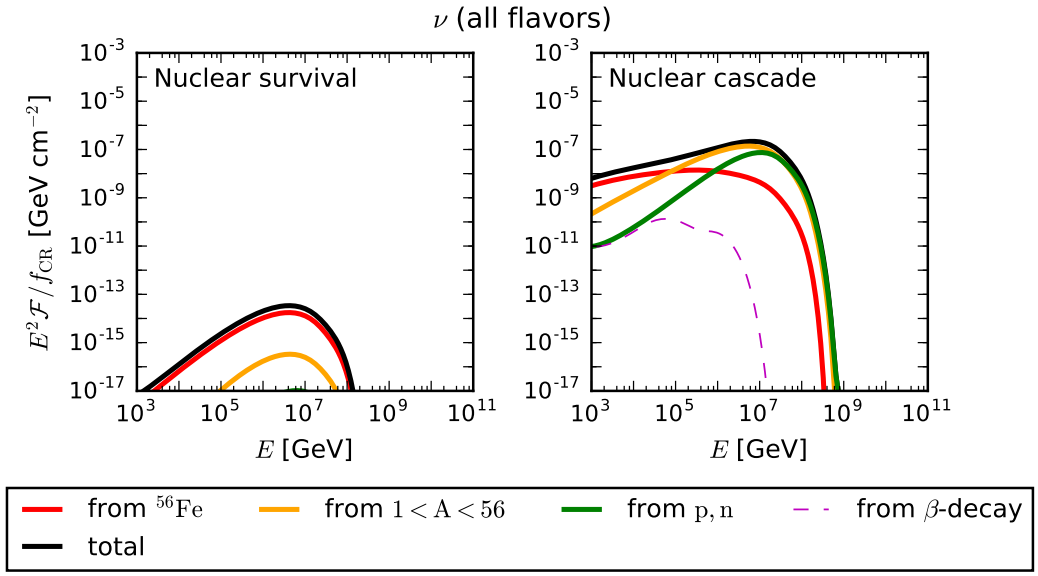

Figure 3.5: Ejected fluence of all-flavor neutrinos Figure 3.2 *(left)* and 3.3 *(right)*, for a pure iron-56 injection composition. We show separately the direct contribution from the primary nuclei, the secondary nuclei heavier than protons, the nucleons, and β decay. The fluences here are shown in the observer's frame, considering a source redshift of $z = 1$ and a CR loading of $\xi_{\rm CR} = 1$ (*cf.* Eq. 2.7). Figure taken from Rodrigues et al. (2018a).

The ejected neutrino fluence for both the nuclear survival and cascade regimes is shown in Figure 3.5, where the contributions from interactions of the primary, the secondary nuclei, the nucleons, and β decay are shown separately (where the secondary nuclei and nucleons are produced in the nuclear cascade). In the nuclear survival regime the low neutrino fluence originates in photo-meson production off the injection isotope (red curve) with very small contributions from secondaries. On the other hand, for sources with a nuclear cascade, most of the neutrinos at high energies are produced off secondary isotopes (yellow curve) and nucleons (green). This is different to the case of GRBs, where there a case is also found where the neutrino production is dominated by emission from secondary nucleons rather than heavier secondaries (Biehl et al., 2017).

As mentioned previously, our simulation of photo-meson production by nuclei is based on a

simplistic superposition model, where the total interaction and inclusive pion production cross sections scale with A (*e.g.* Boncioli et al., 2017). On the other hand, the above results shows that in blazars that are efficient neutrino emitters, a substantial contribution can come from photo-meson production by secondaries heavier than protons, which highlights the importance of more realistic photo-meson production models (see *e.g.* Morejon, 2019).

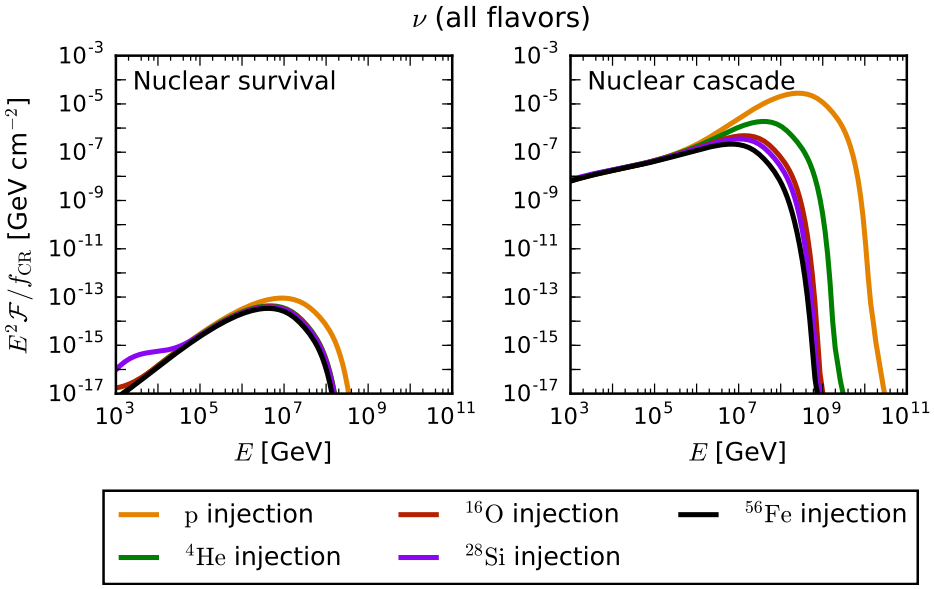

Figure 3.6: All-flavor neutrino fluence for the example of Figure 3.2 *(left)* and 3.3 *(right)*. The neutrino emission is shown for different injected isotopes, in the observer's frame, considering a redshift $z = 1$ and a CR loading of $\xi_{\rm CR} = 1$. Figure taken from Rodrigues et al. (2018a).

Finally, in Figure 3.6 we show the emitted neutrino spectrum for the injection of different isotopes. In the nuclear survival case (left panel), the maximum energy *per nucleon* scales with Z/A (by dividing $E'_{\rm max}$ from Eq. (2.2) by the number of nucleons A). For the most stable isotopes, Z/A has the same value of around 1/2 (except protons), which justifies the weak dependence of the neutrino peak with composition. On the other hand, in the nuclear cascade regime (right panel), the neutrino spectrum basically follows the maximum energy per nucleon as well, which however depends in a non-trivial way on the competition of the disintegration, photo-meson and acceleration time scales. In fact, from Figure 3.4, one can see that the maximum primary energy $E_{\rm max}$ for point B varies only a factor of a few with injection composition, which means that the neutrino energy roughly scales with $E_{\rm max}/A$ following the energy per nucleon. Consequently, the neutrino spectrum extends to higher energies for lighter injection compositions, which also yields a higher neutrino fluence.

3.4 Cosmic ray emission

In Section 3.2 we discussed CR densities inside the jet, relevant for neutrino production as discussed in Section 3.2. In order to address CR emission, we must model the escape mechanism of CRs from the jet.

As discussed in Section 2.1.3, the escape of particles from the source environment on the macroscopic scale is typically modeled via diffusion and advection processes. Here, we consider the two most extreme scenarios introduced in Section 2.1.2: a direct escape by a purely diffusive mechanism, where the escape rate depends on the Larmor radius of the particle at any given energy, and a purely advective mechanism where all particles escape with the speed of light, regardless of energy.

In the upper panels of Figure 3.7 we show the rates for the diffusion-dominated escape as solid curves, and the advective escape rate as a dashed line, corresponding to the free-streaming escape rate, $c/R'_{\rm blob}$. In the nuclear survival case, the maximum energy is limited by the light-crossing time, which implies that even in the diffusive case the Larmor radius reaches the size of the blob at the highest energy, and the escape rate is maximal. In the nuclear cascade case (upper right-hand panel), the maximum energy is limited by photo-disintegration, as shown by the shaded gray area, and at this energy the Larmor radius of the iron nucleus is still around 100 times smaller than the size of the region, due to the strong magnetic field. By Eq. (2.3), the escape will be suppressed by a factor 100 compared to the advective case. While this is true for charged CRs, note that neutrons are also abundantly produced in this case, and being neutral, will not be confined by the magnetic field and will instead escape with the advective timescale regardless of the escape assumption.

The second and third rows in Figure 3.7 compare the effect of the two escape hypotheses on the ejected CR spectra. The total escape spectra in the advective case resemble the E^{-2} injection spectrum over a wide range of energies, since the escape rate is flat in energy, *i.e.* they are generally softer than in the diffusive case[1]. The diffusion assumption, on the other hand, hardens the spectrum of the escaping charged particles (*cf.* Eq. 2.3 and discussion thereafter). As mentioned earlier, the neutron spectra are independent of the escape mechanism, since they are not confined by the magnetic fields in the source.

Finally, in the fourth row of Figure 3.7 we show the average (logarithmic) mass of ejected CRs, $\langle \ln A \rangle$. Unsurprisingly, in the nuclear survival case (left), the emitted composition is the same as that of the injected CRs. A more interesting case is the advection scenario (right), where diffusion and advection yield opposite trends of the ejected CR composition. Drawing

[1] Note that in the nuclear survival case (left-hand side panels), near the spectral cutoff, the fluence is higher for the advective escape assumption simply because of the absence of an adiabatic cooling term.

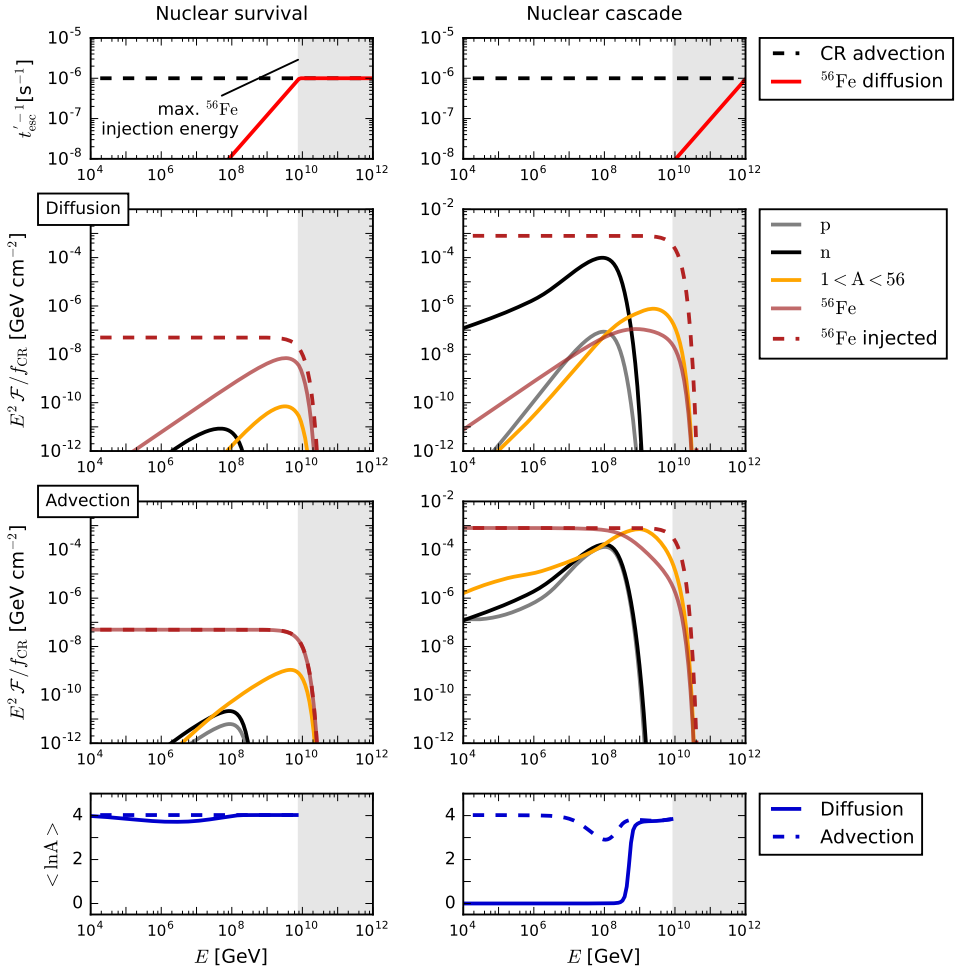

Figure 3.7: Cosmic ray escape rates, ejected CR spectra and average CR composition for the nuclear survival and nuclear cascade examples. *Top row:* Escape rates (in the blob rest frame) for a diffusion (solid) and advection (dashed) escape mechanism. *Second row:* Ejected spectra for diffusive CR escape, given in the observer's frame considering a redshift $z = 1$. *Third row:* Ejected spectra for advective CR escape. *Bottom row:* Average ejected CR composition for both escape mechanisms. The energies on the horizontal axes are in the observer's frame. Note that no interactions in the propagation are included. Figure taken from Rodrigues et al. (2018a).

from the discussion in Section 2.1.1, a diffusive escape mechanism can potentially lead to a trend of the ejection composition as a function of energy that is more compatible with observations (Aab et al., 2017b, *cf.* Figure 2.1). Note, however, that CR propagation also plays a role in the observed composition, and this has not been taken into account in this study.

3.5 A three-zone model for FSRQs

As discussed in Section 2.4.4, FSRQs typically have a large BLR and bright disk emission, and external fields like broad line and thermal emission can therefore be transported into the jet. Assuming the relationships given by Eqs. 2.15 and 2.17, only high-luminosity FSRQs (HL-FSRQs) may have BLR and dust torus radiation zones large enough to physically contain the blob (as in the middle panel of Figure 2.8), yielding an external-field scenario. With the baseline parameters considered in this chapter, this happens in blazars with disk luminosity $L_{\rm disk} \gtrsim 10^{46}$ erg/s. The spectrum of the external fields follows Murase et al. (2014), and an example was given in Figure 2.9, which corresponds to the same nuclear cascade prototype discussed in this chapter. For different luminosities, the external fields simply scale correspondingly in normalization without any change in their spectrum.

Besides interacting with the external fields inside the blob, it is also important to realize that after escaping the jet, CRs must travel large distances inside these external radiation fields before escaping into the intergalactic medium. This implies additional CR cooling and neutrino emission that is not accounted for in a jet-only model. On the contrary, for less powerful FSRQs and also for BL Lacs, a single plasma blob scenario containing only non-thermal jet emission characterizes well the interaction environment. As shown previously in Figure 2.9, the density of these external fields is much lower in the black hole frame than in the shock rest frame of the jet, since in the jet they appear Lorentz-boosted according to Eq. (2.18). However, note that the CRs escaping the jet carry the Lorentz factor of the jet itself, and therefore in the rest frame of the CRs the external fields still appear Lorentz-boosted, like in the right panel of Figure 2.9.

In Figure 3.8 we show the interaction rates in the jet blob, the BLR and the dust torus radiation field, for the same blazar shown previously as an example of the nuclear cascade scenario. In the left panel we can see that the external fields boost photo-disintegration in the jet, which becomes optically thick above PeV energies (compare dotted and solid green curves). Neutrino production from photo-meson interactions of PeV protons also becomes more efficient by two orders of magnitude (compare dotted and solid red curves). The nuclei that escape from the jet zone into the BLR and dust torus then undergo additional disintegration above 10 PeV, thus reducing the flux of heavy UHECRs emitted by the source. Finally, we see that the dust torus only affects hadronic interactions of CRs above a few EeV (green curves in the middle and

right panels of Figure 3.8).

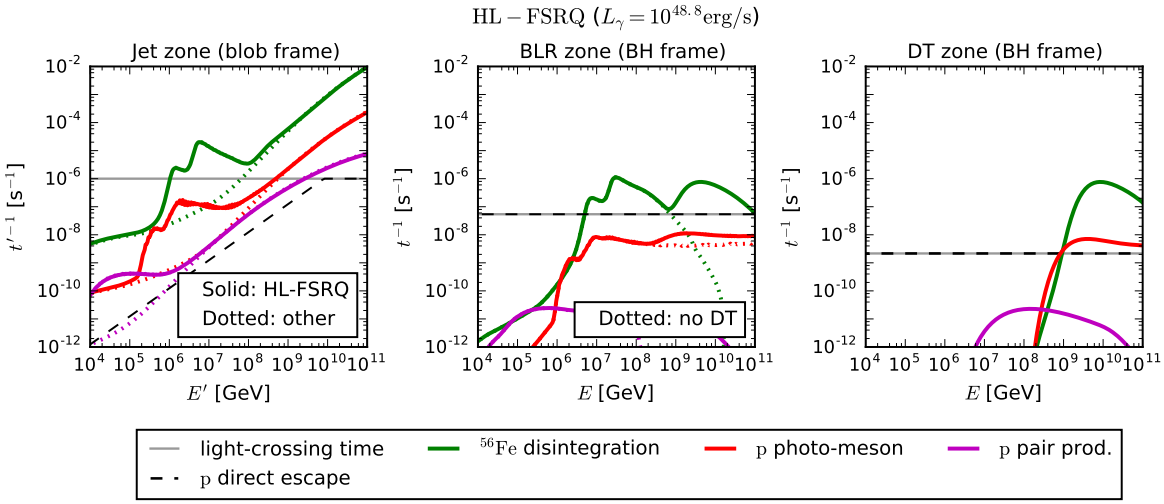

Figure 3.8: Interaction rates in the three radiation zones of the FSRQ model, in the rest frame of the zone, for the same source as in Figure 2.9. By comparing the solid and dotted lines, we see the effect of the external radiation fields in the inner zones of the model. Figure taken from Rodrigues et al. (2018a).

We assume the magnetic field in the BLR and dust torus to be much weaker than in the jet, $B_{\rm BLR} \ll 1$ mG, a valid assumption if these regions consist of gas and clouds similar to supernova remnants. This implies that no synchrotron losses and no magnetic confinement takes place in the external zones.

In practice, the three-zone model is computed numerically by re-injecting the CRs that escape from the jet into the BLR, and then subsequently into the dust torus radiation field (*cf.* Figure 2.8). The details of this procedure are presented in Appendix B. Essentially, the solver first integrates the blob system up to a few dozen times $t'_{\rm flare}$, until all the CRs that were accelerated during the duration of the flare have either cooled or escaped the blob. The CRs that escape from the blob are then injected into a new set of equations describing the radiation environment of the BLR, and after the BLR the same process takes place for the dust torus zone.

The effect of the CR propagation in the external radiation fields in the emitted neutrino spectra are shown in Figure 3.9 for the two escape assumptions considered in Section 3.4. Here, the neutrino spectra emitted from CR interactions in each zone are shown separately. The advection case yields in general more neutrinos because there are more CRs escaping the jet into the BLR, is populated with a higher initial density at $t' = t'_{\rm flare}$. Therefore, a higher number of primaries is available for neutrino production in the outer zones compared to the diffusion

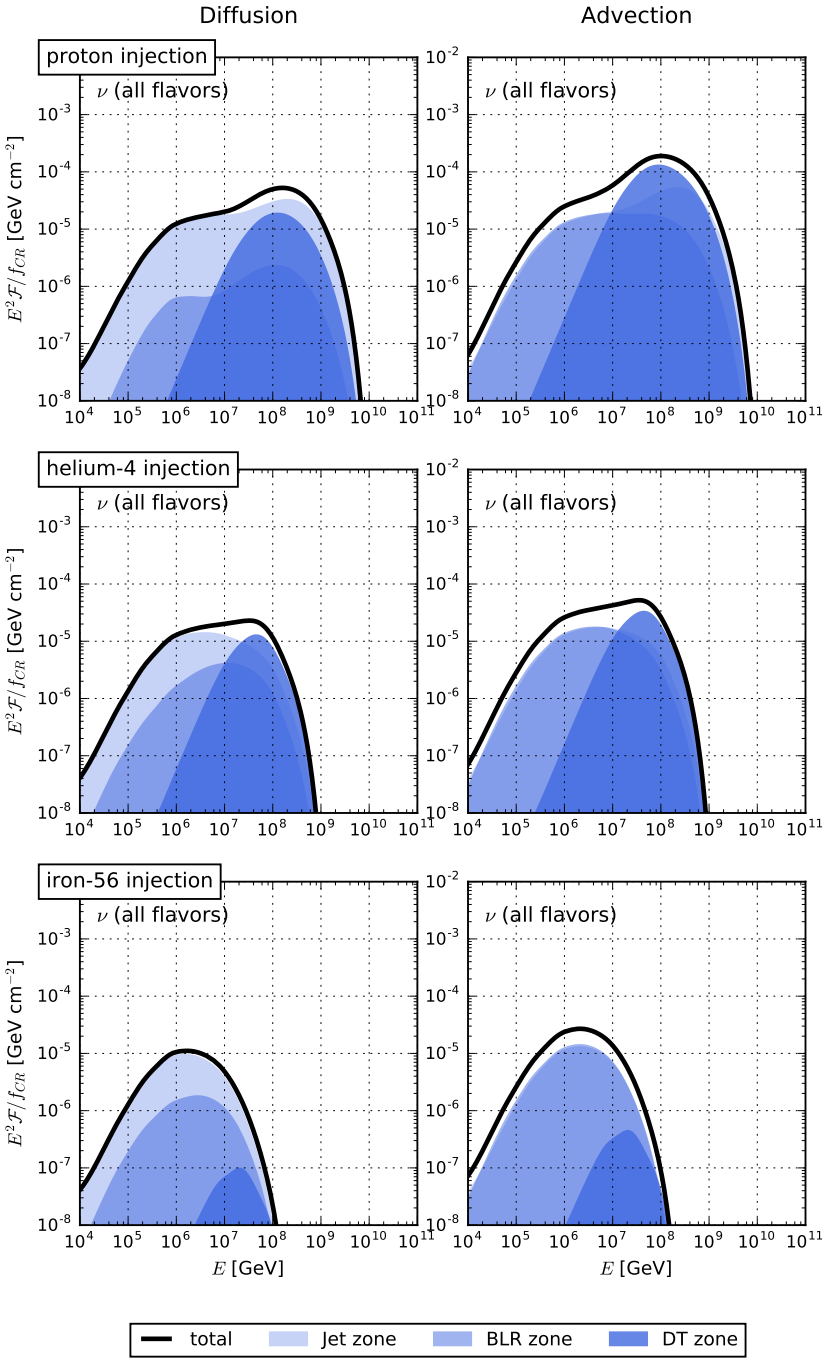

Figure 3.9: Ejected all-flavor neutrino fluence (in the observer's frame) for the injection of a pure composition of protons *(top)*, helium-4 *(middle)* and iron-56 *(bottom)*, for an FSRQ with $L_\gamma = 10^{48.8}$ erg s^{-1} (see Tab. A.1) at $z = 1$, assuming the CR escape mechanism to be diffusive *(left)* and advective *(right)*. The contributions from the different regions are plotted separately. Note that the fluences from the three blazar zones add up to the total ejected spectrum, since neutrinos free-stream out of the source. Figure taken from Rodrigues et al. (2018a).

case (darker blue shaded regions Figure 3.9). For the pure proton injection (upper right panel) and advective escape, the outer zones (in particular the DT) notably broaden the peak at tens of PeV and increase the flux by factor ~ 2 compared to diffusive escape. This result is found to be consistent with previous studies (Murase et al., 2014).

The injection of helium-4 (middle panels) results in a similar behavior, except that the jet contribution is smaller as already discussed in Section 2.2. For the same reason as in the proton case, the advection case produces more neutrinos. Since the BLR is optically thick for heavier nuclei, most of the secondaries from iron-56 injection (lower panels) are absorbed in the BLR, such that neutrino emission from the dust torus is lower compared to the other two cases. Since HL-FSRQs clearly lie inside the nuclear cascade regime, they emit mostly nucleons and photo-disintegrated secondary nuclei.

3.6 Application to the blazar sequence

In this section, we apply the hadronic model with nuclei to blazars of different luminosities, using the concept of blazar sequence introduced in Section 2.3.2. Despite the debate on whether this is influenced by a selection bias due to the limited sensitivity of gamma-ray telescopes, this phenomenological prescription does adequately describe at least the sample of the observed blazars, providing an average SEDs in each luminosity bin. These SEDs are given in Figure 3.10 in the rest frame of the jet, and the parameters of the model are detailed in Tab. A.1, in Appendix A. As explained previously, we only consider the effect of external fields for high-luminosity blazars, which correspond to the three brightest SEDs of Figure 3.10. For these sources, we apply the three-zone model introduced in the last section, while for dimmer objects a one-zone model is used. To study the case of low-luminosity BL Lacs, we extrapolate the sequence down to $L_\gamma \sim 10^{42}$ erg s^{-1} (dashed SEDs in Figure 3.10).

Figure 3.11 shows the results for the maximum CR energy ejected by each source, for different injected isotopes and for either of the escape assumptions discussed in Section 3.4. For HBLs, where photo-nuclear/-hadronic interactions are sub-dominant up to the highest injection energies, the maximum energy is determined from the Hillas condition in the blob. As a result, we have $E_{\rm max} \propto B' \propto L_\gamma^{1/2}$ along the blazar sequence, since a constant energy partition $\epsilon_{\rm B}$ is assumed for the magnetic fields. In the range indicated by the gray band, the maximum energy is proportional to the charge of the injected isotope, which corresponds to the requirement that cosmic accelerators follow the Peters cycle (Peters, 1961), as introduced in Section 2.1.2. This relationship breaks down at the transition between the nuclear survival and nuclear cascade regimes at $L_\gamma \gtrsim 10^{45} - 10^{46}$ erg s^{-1}, where photo-hadronic interactions become efficient and limit the maximum energy of ejected CRs. A further increase of the luminosity results

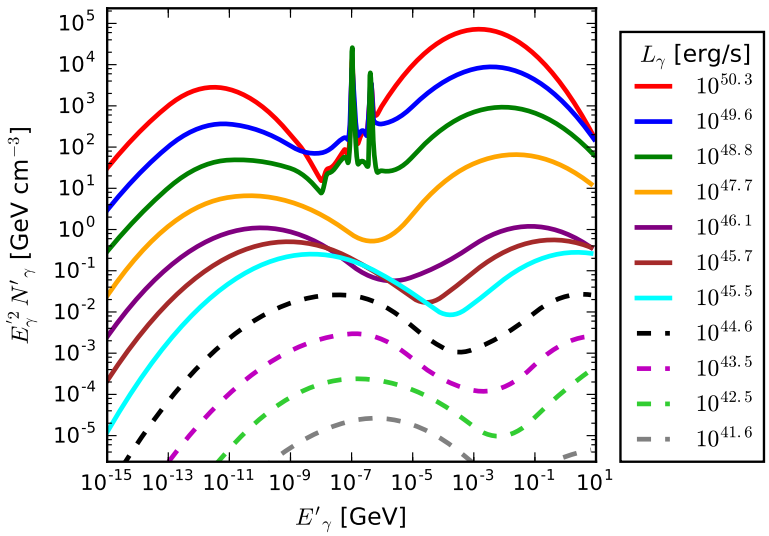

Figure 3.10: Representation of the average SEDs of blazars of different luminosities, given in the reference frame of the jet, following the original blazar sequence by Fossati et al. (1998). The solid lines are taken from Murase et al. (2014), while the dashed lines are extrapolated SEDs representing the extension of the sequence to high-synchrotron-peaked BL Lacs. Note that the contribution from external fields in the three brightest SEDs (the two broad lines and the thermal continuum from the accretion disk) appears boosted from within the jet but not in the observer's frame, where only the two broad bands from non-thermal emission can be detected. Figure taken from Rodrigues et al. (2018a).

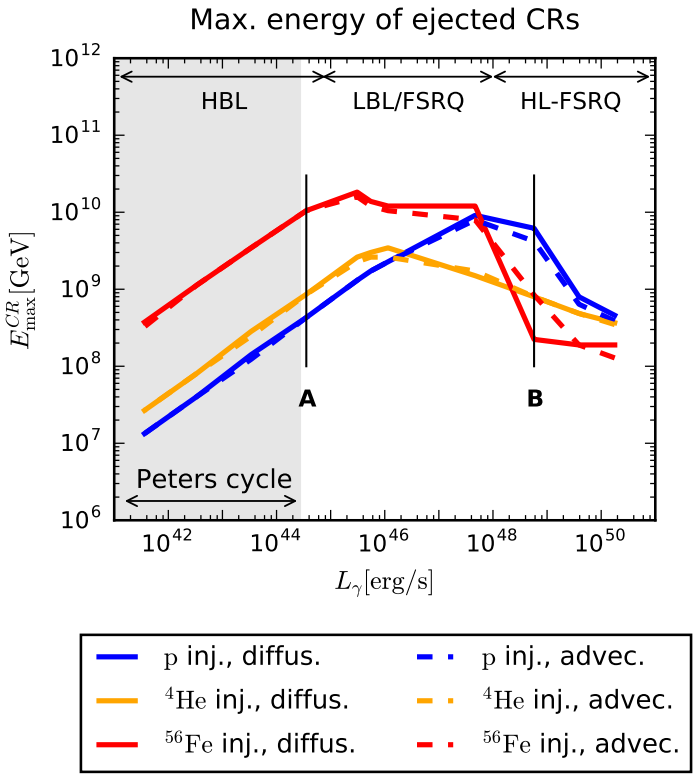

Figure 3.11: Maximum ejected CR energy for each source described in Tab. A.1 and Figure 3.10. Only BL Lacs emit CRs that follow the Peters cycle $E_{\rm max} \sim Z$, *i.e.* a scaling of the maximum energy with Z. The luminosity marked A corresponds to the nuclear survival prototype from Section 3.2, while B corresponds to the HL-FSRQ explored in Figures 2.9 to 3.9. Figure taken from Rodrigues et al. (2018a).

in a reduction of the maximum energy for the reasons discussed in Section 3.2. The effect is more pronounced for heavier injection isotopes, as we can see from the three cases displayed in Figure 3.11. Note that the maximum energy is not significantly affected by the escape scenario.

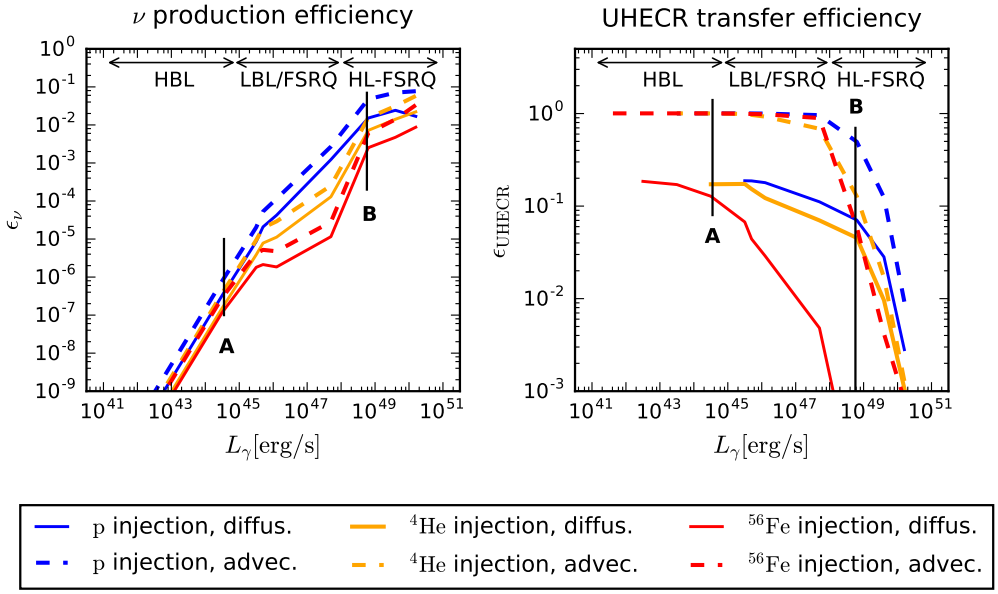

Figure 3.12: *Left:* All-flavor neutrino production efficiency for blazars of different luminosities, following the blazar sequence described in Tab. A.1 and Figure 3.10. *Right:* Cosmic-ray energy transfer efficiency in the UHE range ($E_{\rm CR} > 10^9$ GeV in the black hole frame) for the same sources. The colors represent three different pure injection compositions; solid and dashed curves represent diffusive and advective escape, respectively. The luminosity marked *A* corresponds to the nuclear survival prototype from Section 3.2, while *B* corresponds to the HL-FSRQ explored in Figures 2.9 to 3.9. Figure taken from Rodrigues et al. (2018a).

We now turn to the production efficiency of UHECRs and neutrinos across the blazar sequence. We define the *UHECR transfer efficiency* of a blazar as the ratio of the total energy emitted during the flare in UHECRs (above 1 EeV in the black hole frame), and the total injected energy in CRs above the same energy, during the dame period.

$$\epsilon_{\rm CR} = \frac{\sum\limits_{\rm allCRs} \int_0^\infty \int\limits_{1\ \rm EeV} \left(\frac{dN}{dE'dt'}\right)_{\rm esc} dE'\,dt'}{\int_0^\infty \int\limits_{1\ \rm EeV} \left(\frac{dN}{dE'dt'}\right)_{\rm inj} dE'\,dt'}\,. \tag{3.3}$$

Similarly, the *neutrino production efficiency* ϵ_ν is obtained by performing the sum in the numerator over all neutrino species, and with both integrals over all energies. Neither of these two

quantities depends on the CR loading factor $\xi_{\rm CR}$, which cancels out in Eq. (3.3) (and also in the formula for ϵ_ν, since the total neutrino fluence is directly proportional to the injected CR fluence).

The neutrino production efficiencies are shown in the left panel of Figure 3.12. Clearly, HL-FSRQs can be identified to be the best neutrino emitters of the blazar sequence, with $\epsilon_\nu \sim 10^{-3}$ to 10^{-1}. For less luminous blazars the photon field densities are too low for efficient neutrino production, with HBL Lacs being the least efficient neutrino sources. From the same plot we can see that HBL Lacs follow the expected relation $L_\nu \sim \epsilon_\nu L_\gamma \sim L_\gamma^2$; above $L_\gamma = 10^{45}$ erg s^{-1}, this scaling becomes weaker, and beyond $L_\gamma = 10^{49}$erg/s it becomes $L_\nu \sim L_\gamma$. There is hardly any dependence on the UHECR escape mechanism, apart from the ~ 2 times more efficient production for the advective case, as discussed in Section 3.5.

On the right-hand side panel of Figure 3.12 we plot the UHECR transfer efficiency. The overall tendency is the opposite of the neutrino efficiency: HL-FSRQs cannot efficiently emit UHECRs due to strong photo-hadronic interactions and also to the additional radiation zones that suppress the UHECR part of the spectrum, while lower-luminosity objects such as LL-FSRQs and BL Lacs may allow for more efficient CR survival and escape. HBL Lacs are expected to be efficient sources of UHECR regardless of the injected isotope, since 100% of the injected UHECRs are emitted. In the advective escape scenario (dashed curves in Figure 3.12) blazars become efficient UHECR source candidates. Note, however, that the spectra of emitted CR are generally softer. For diffusion-dominated CR escape, the production efficiency is low. The UHECR transfer efficiency should be interpreted together with the maximum achievable energy in (see Figure 3.11): While, for example, at low energies a large fraction of the injected energy is transferred, the magnetic field is too low to allow for very high UHECR energies of light isotopes. Therefore, each injection isotope has a "sweet spot" where UHECR can be accelerated to sufficiently high energies and escape from the source environment. For example, for helium injection, this sweet spot seems to be around $10^{46}\,{\rm erg\,s^{-1}}$. If, in addition, neutrinos are to be efficiently produced, higher luminosities are preferred; $L \sim 10^{48}\,{\rm erg\,s^{-1}}$ seems to be a good trade-off between neutrino and UHECR production.

3.7 Conclusion

In this chapter we studied blazars as sources of UHECRs and neutrinos including the injection of isotopes heavier than hydrogen in the jet. We can identify two distinct regimes, depending on jet luminosity and size of the blob: in the nuclear survival regime, corresponding to low luminosities or large blob sizes, the source is optically thin to photo-nuclear interactions and neutrinos are mostly produced by photo-meson production off the primary nucleus and the maximum CR

energies behave as expected from the Peters cycle, *i.e.* $E_{\rm max} \sim Z$. In the nuclear cascade regime, corresponding to high luminosities or small blob sizes, the source is optically thick to disintegration of the injected nuclei at the highest energies and a nuclear cascade develops in the jet. Neutrinos are produced in high amounts by photo-meson interactions of the secondary nuclei generated in the nuclear cascade, and the maximum CR energies saturate due to photo-disintegration of UHE nuclei.

The evolution of neutrino and CR production efficiencies has been studied over the blazar sequence. In order to have a more refined model for HL-FSRQs, we have included the external radiation fields in the jet and the propagation of CRs through the radiation fields of the BLR and dust torus. The neutrino production efficiency strongly increases with luminosity from a low value for HBLs, over LBLs/FSRQs, and then saturates for HL-FSRQs. For HBLs we have $L_\nu \propto L_\gamma^2$, whereas for HL-FSRQs we find $L_\nu \propto L_\gamma$ (Figure 4.3). On the contrary, the UHECR transfer efficiency is high for HBLs, then decreases for LBLs/FSRQs, and it is low for HL-FSRQs. From Figure 4.3 we also deduce that here may be a "sweet spot" in the range around $L_\gamma \sim 10^{48}$ erg s^{-1} where both UHECRs and neutrinos are efficiently produced.

We also discussed the impact of the UHECR escape mechanism, focusing on a diffusion scenario (leading to hard ejection spectra), and an advection scenario (leading to an unmodified spectrum compared to the acceleration spectrum, except at the highest energies where hadronic cooling takes place). Both spectrum and composition of the diffusion scenario agree better with recent Auger observations (Aab et al., 2017b), although propagation effects can change the composition, and different acceleration mechanisms may accelerate CRs to spectra harder than E^{-2}. The neutrino production in the jet is not affected by the escape assumption, but the external zones (BLR and dust torus) should play a more important role in the case of advective escape.

We conclude that HL-FSRQs are very efficient neutrino emitters, whereas UHECRs should come preferentially from low- and intermediate-luminosity blazars. These conclusions pertain to individual sources, while the general picture must include a model of the cosmological blazar distribution and how some key physical properties may change among blazars. For example, in spite of their low neutrino luminosity, a sufficiently large number of LBLs that are not detectable in gamma-rays can conceivably dominate the diffuse neutrino flux. Another variable is the CR loading, which may vary for different blazar families or evolve over the blazar sequence. In the next chapter, we will apply this model to an entire blazar population in order to study this issue.

Chapter 4

Neutrino emission from the entire blazar family

In this chapter we study the contribution of blazars to the astrophysical neutrino flux. As discussed in Section 2.2, considering the high-energy starting events (HESE) in IceCube, the contribution of resolved blazars to the astrophysical neutrino flux should not be larger than $\sim 20\% - 25\%$ (Aartsen et al., 2017a) based on the missing association with sources in gamma-ray catalogs, or even less based on theoretical considerations (Palladino & Vissani, 2017). On the other hand, there have been indications of associations of individual neutrino events with AGNs (Resconi et al., 2017; Padovani et al., 2016; Kadler et al., 2016; Aartsen et al., 2018b), which may indicate a partial contribution from blazars to the diffuse neutrino flux.

The neutrino emission efficiency of each individual source, ϵ_ν, is obtained using the methods explained in the previous chapter, but considering different source parameters such as magnetic field strength and acceleration efficiency η. In Chapter 3 the acceleration efficiency was considered tho have the maximum value of $\eta = 1$. Such ultra-efficient acceleration leads to the ejection of UHECRs up to a few EeV (Figure 3.7), and would therefore be necessary if we intended to explain the high-energy tail of the observed CR spectrum; however, the neutrinos emitted would have equivalently high energies of order $\sim$ EeV (Figure 3.6), which is not suitable for explaining the currently measured flux of astrophysical neutrinos, which extends up to only a few PeV (Figure 2.2). Moreover, in this chapter the baryonic loading $\xi_{\rm CR}$ (see Eq. 2.7) will be allowed to vary for blazars of different luminosities. We will show that low-luminosity blazars can provide the dominant contribution to the high-energy neutrinos with energy between a few hundreds of TeV and a few PeV, while the contribution of very bright sources to the neutrino flux must be highly suppressed in order to respect the blazar stacking limit (Aartsen et al., 2017a). We also discuss explicitly the role of TXS 0506+056, and what we can learn from the population study for future such observations.

The study presented in this chapter is based on Palladino et al. (2019). My main contribution to this project was the calculation of the neutrino spectrum and neutrino efficiency of blazars of different luminosities, using the hadronic blazar model presented in the previous chapter.

4.1 Blazars as sources of the observed high-energy astrophysical neutrinos

The calculation of neutrino production in blazars follows the same methods used in the previous chapter. However, we re-consider the following model assumptions that differ from those in the previous chapter:

- The most recent version of the blazar sequence (Ghisellini et al., 2017) is adopted, based on the *Fermi* 3LAC catalog, as opposed to the old sequence by Fossati et al. (1998), which was based only on a few sources. Moreover, in the previous chapter we neglected low-luminosity FSRQs and high-luminosity BL Lacs, due to the reduced number of both these source types. On he contrary, here we consider the full sequence for BL Lacs and FSRQs separately.

- We adopt a value of acceleration efficiency $\eta = 10^{-3}$ (*cf.* Eq. 2.1), compared to the value of $\eta = 1$ used in the previous chapter. This allows for an emitted neutrino spectrum that peaks at energies compatible with the IceCube diffuse flux (Figure 2.2) rather than EeV energies like in the previous chapter. While the motivation for this choice is phenomenologically motivated, there is evidence that relativistic shocks may in fact be inefficient accelerators through diffusive shock acceleration (Inoue & Tanaka, 2016).

- The magnetic field strength is assumed to scale as a soft power law of the gamma-ray luminosity, yielding 0.1 G for the dimmest BL Lac considered, and 6 G for the brightest FSRQ; these magnetic field strengths are typical of blazars of the respective luminosities.

- The effect of the dust torus radiation field is taken into account for intermediate-luminosity FSRQs, where the jet blob is located outside the BLR but within the dust torus radius, $R_{\rm DT} > R_{\rm diss} > R_{\rm BLR}$ (*cf.* Section 2.4.4).

- In the previous chapter, we modeled the scenarios of a purely advective and a purely diffusive CR escape mechanism (see also Section 2.1.3). Here, we limit ourselves to the purely diffusive escape case, which means we assume a rigidity-dependent CR escape rate. This is justified by the fact that in this model, the neutrino production efficiency does not strongly depend on the escape mechanism (see Figure 3.12).

- In this study only protons are considered. The reason for this choice is two-fold: firstly, a population study will unavoidably have a great number of parameters, and therefore the addition of chemical composition would imply an additional degree of complexity to what is already an intricate problem. Secondly, as we learned from Figure 3.12, the *scaling* of

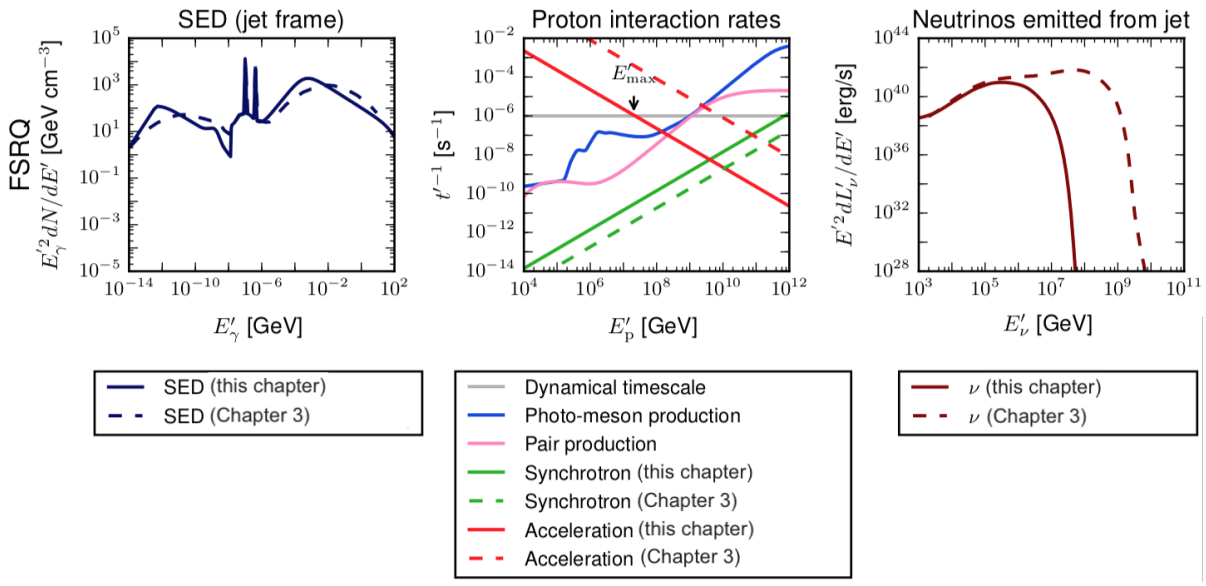

Figure 4.1: *Left:* the solid curves represent the photon density spectra in the jet in an FSRQ of luminosity $10^{48.5}$ erg s^{-1} from the updated blazar sequence, considered in this chapter. For comparison, we plot in dashed the SED of the FSRQ used as an example in the previous chapter (Figure 3.3). *Middle:* cooling and interaction rates of protons in the jet for the same sources. The solid red line is the acceleration efficiency considered in this chapter, $\eta = 10^{-3}$. For comparison, the red dashed line refers to an ultra-efficient proton acceleration ($\eta = 1$). The stronger synchrotron cooling rates (compare solid and dashed yellow lines)are due to the higher magnetic field strength, following the discussion in the main text. *Right:* all-flavor neutrino luminosity spectrum produced in the jet, considering the SED and acceleration efficiency used in this chapter (solid) compared to the discussion in Chapter 3 (dashed). Figure adapted from Palladino et al. (2019).

the neutrino production efficiency ϵ_ν with the blazar luminosity (which in this work is the only information we wish to obtain from the source model) does not depend dramatically on the injected chemical composition, particularly at low luminosities.

Other model parameters are kept equal to the previous chapter, such as the Lorentz factor $\Gamma = 10$ and the blob radius $R'_{\rm blob} = 3 \times 10^{16}$ cm.

The effect of the changes enumerated above are summarized in Figure 4.1. On the left panel we see to SEDs of the same gamma-ray luminosity, one from the old blazar sequence, considered in the previous chapter (dashed) and from the updated blazar sequence, as considered in the present chapter (solid). The difference relates only to the non-thermal SED, while the features from external fields (accretion disk, dust torus and broad lines) are the same (*cf.* Figure 2.9).

In the middle panel we plot the interaction rates obtained with the current SED (solid) and we also show a comparison of the acceleration and synchrotron rates with the previous chapter. The higher synchrotron rates (green lines) are due to the different scaling of the magnetic field with gamma-ray luminosity, as explained previously, which means that a blazar with this particular gamma-ray luminosity will now have a lower magnetic field strength in the jet. On the other hand, the acceleration rates (red lines) are lower for this case due to the lower value of acceleration efficiency. Finally, on the right plot of Figure 4.1 we see a comparison of the ejected neutrino spectrum with each of the parameter sets (the dashed curve corresponds to the spectrum shown in light blue in the upper right-hand panel of Figure 3.9). The difference in the peak and total fluence of the emitted neutrinos is due to the lower acceleration efficiency considered in the resent case, which yields a lower maximum energy of protons accelerated in the source.

The SEDs considered for different luminosities are represented in the left panels of Figure 4.2, in the rest frame of the jet. The neutrino spectra produced as a result of the model are shown in the right-hand panels of the same figure. The integral of the neutrino spectrum of each blazar yields the total neutrino luminosity of the source, L_ν. This quantity is directly proportional to the baryonic loading $\xi_{\rm CR}$ of the source (see Eq. 2.7), since the emitted neutrinos originate from CR interactions (Figure 4.2 refers to the special case $\xi_{\rm CR} = 1$). In general, the emitted neutrino luminosity L_ν of a blazar may be written as

$$L_\nu = \mathcal{K}\, L_\gamma, \tag{4.1}$$

where $\mathcal{K} \equiv \epsilon_\nu \times \xi_{\rm CR}$ is the product of the neutrino production efficiency and the baryonic loading, defined in Eq. (2.7). As discussed in the previous chapter, the neutrino production efficiency $\epsilon_\nu \equiv L_\nu / L_{\rm CR}$ quantifies the energy from CRs converted into neutrinos in the source, and is independent of the baryonic loading. Throughout this chapter the quantity $\mathcal{K} = L_\nu / L_\gamma$ will be used to express the ratio between the neutrino and gamma-ray luminosity of a given blazar. While a number of works limit the discussion to the case $\mathcal{K} = $ const. (*e.g.* Wang & Li, 2016; Kadler et al., 2016; Righi et al., 2017; Halzen & Kheirandish, 2016), in this study we will also explore scenarios where $\mathcal{K}$ varies as a function of the blazar luminosity.

Using the results of Figure 4.2, we can calculate the neutrino production efficiency of each source considered in this work. This is shown in the top panel of Figure 4.3 as a function of the gamma-ray luminosity of the source. The neutrino production efficiency increases monotonically with the gamma-ray luminosity of the blazar, since higher luminosities imply higher photon densities in the radiation zone in the jet. Note the abrupt increase in the efficiency of FSRQs with $L_\gamma \gtrsim 3 \times 10^{48}$ erg/s, which is due to interactions with external fields, producing additional

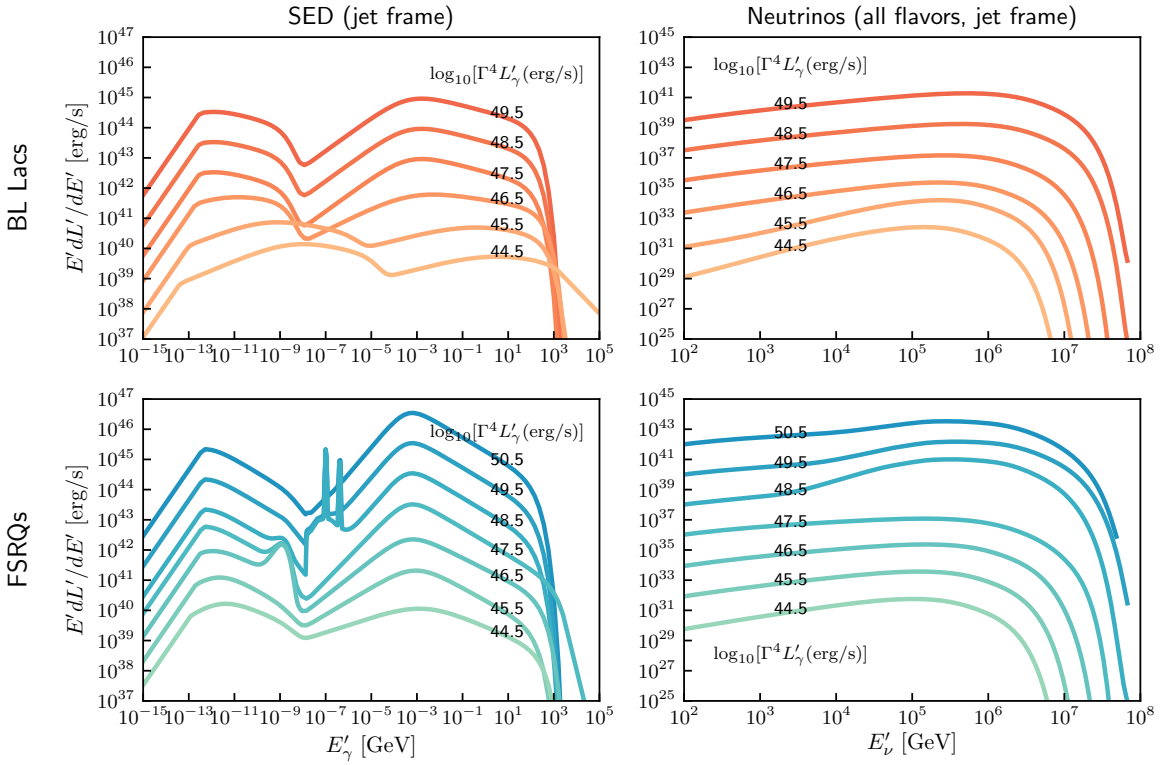

Figure 4.2: *Left:* SEDs for BL Lacs (top) and FSRQs (bottom) used in this work, given in the rest frame of the jet. The non-thermal SEDs are those from the blazar sequence (Ghisellini et al., 2017), and the external components of the energy spectra are based on Murase et al. (2014). *Right:* neutrino luminosity spectra calculated for each blazar luminosity, for a baryonic loading of $\xi_{\mathrm{CR}} = 1$, also in the rest frame of the jet. The label of each curve indicates the logarithm of the gamma-ray luminosity in the black hole frame in erg/s, assuming a bulk Lorentz factor of the jet of $\Gamma_{\mathrm{b}} = 10$. Figure taken from Palladino et al. (2019).

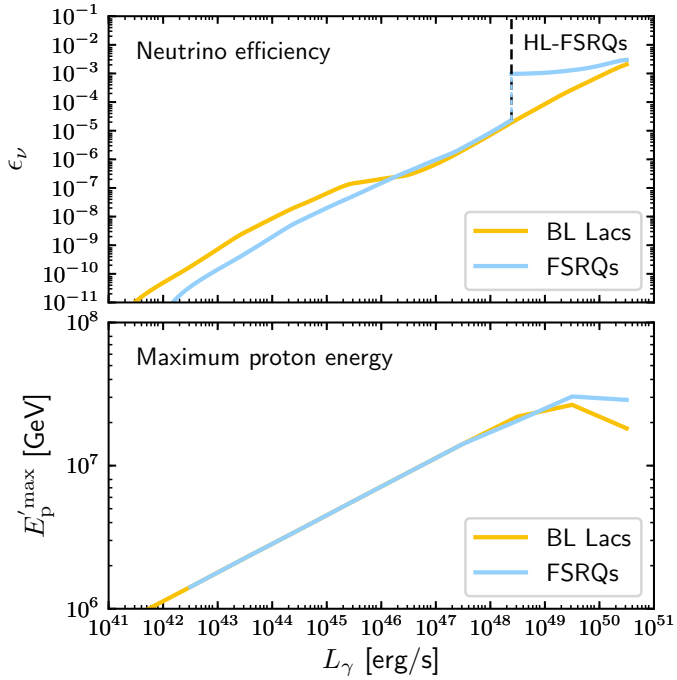

Figure 4.3: *Top:* neutrino production efficiency of BL Lacs (yellow) and FSRQs (blue) of the blazar sequence as a function of the gamma-ray luminosity of the source. *Bottom:* maximum energy of accelerated protons in the jet of BL Lacs (yellow) and FSRQs (blue) as a function of luminosity. We assume that the radiation zone in the jet is exposed to external photon fields from molecular and thermal emission in the case of FSRQs with luminosity above 3×10^{48} erg/s (HL-FSRQs); this is responsible for the jump in neutrino efficiency. Figure taken from Palladino et al. (2019).

neutrinos.

In the bottom panel of Figure 4.3 we show the maximum energy achieved by protons accelerated in the jet, which as we can see is highest for blazars of $L_\gamma \sim 3 \times 10^{49}$ erg/s. Above this luminosity, photo-hadronic interactions of high-energy CRs dominate over acceleration, and the maximum CR energy starts decreasing with luminosity.

4.1.1 The expected flux of astrophysical neutrinos

The description of the blazar population is based on the distribution of BL Lacs and FSRQs by Ajello et al. (2014, 2012), which was introduced in Section 2.3.2. In these two works the distribution of blazars is parametrized in order to reproduce the sources observed by Fermi-LAT and the observed diffuse gamma-ray background (Abdo et al., 2010b). The complexity of this

model of BL Lacs and FSRQ evolution is in contrast with other simplified assumptions adopted in other population studies in the literature (see *e.g.* Gelmini et al., 2012, who assume a simple scaling of $(1+z)^5$ for all AGNs).

The expected flux of astrophysical neutrinos from blazars can be determined using the neutrino flux from each source (given its luminosity and redshift) and the cosmological distribution of blazars. The neutrino flux at Earth dF_ν/dE produced by a single source with gamma-ray luminosity L_γ at redshift z is given by the following expression:

$$\frac{dF_\nu}{dE}(L_\gamma, z, E, \xi_{\rm CR}(L_\gamma)) \;=\; \frac{1}{4\pi D_c^2(z)}\frac{dL_\nu}{dE}(L_\gamma, (1+z)E), \tag{4.2}$$

where dL_ν/dE is the neutrino luminosity spectrum of each source, represented in the right panels of Figure 3.5, and D_c is the comoving distance to the source. The relationship between the total neutrino luminosity of the blazar and its gamma-ray luminosity is given by Eq. (4.1), and the respective neutrino efficiency values ϵ_ν are given in Figure 4.3.

The flux of neutrinos observed at Earth form an entire blazar population can therefore be obtained from integrating the neutrino flux of each source (Figure 4.3) with the source's luminosity function ((represented in Figure 2.6), following Ajello et al., 2012, 2014). This is represented by the expression:

$$\begin{aligned}\frac{d\Phi_{\rm tot}}{dE}(E, \xi_{\rm CR}(L_\gamma)) = &\left[\int_{z_1}^{z_2}\int_{L_{\gamma 1}}^{L_{\gamma 2}} \frac{dN}{dz dL_\gamma} \times \frac{dF_\nu}{dE}(L_\gamma, z, E, \xi_{\rm CR}(L_\gamma))\; dz dL_\gamma\right]_{\rm BLLacs} \\ &+ \left[\int_{z_1}^{z_2}\int_{L_{\gamma 1}}^{L_{\gamma 2}} \frac{dN}{dz dL_\gamma} \times \frac{dF_\nu}{dE}(L_\gamma, z, E, \xi_{\rm CR}(L_\gamma))\; dz dL_\gamma\right]_{\rm FSRQs}.\end{aligned} \tag{4.3}$$

In general, the integration limits $L_{\gamma 1}$, $L_{\gamma 2}$, z_1 and z_2 should be such that the entire source distribution is accounted for (right panel of Figure 2.6). If, however, we wish to obtain separately the contributions of unresolved and resolved blazars, the integrals in Eq. (4.3) must be performed only in the range where the sources are detectable by *Fermi*. As discussed in Section 2.3.2, this means that the integration in luminosity must be performed only below $L_\gamma^{\rm vis}(z)$ for unresolved sources, and only above $L_\gamma^{\rm vis}(z)$ for resolved sources, with $L_\gamma^{\rm vis}(z)$ being the luminosity of a blazar at a given redshift that yields a flux detectable bu *Fermi*, $\phi_\gamma^{\rm vis} = 4\times10^{-12}$ erg cm^{-2} s^{-1}. This translates into the relationship $L_\gamma^{\rm vis}(z) = 4\pi \times \phi_\gamma^{\rm vis} \times D_L^2$, where D_L is the luminosity distance of the source. Note that with this treatment we are assuming a constant average flux from blazars.

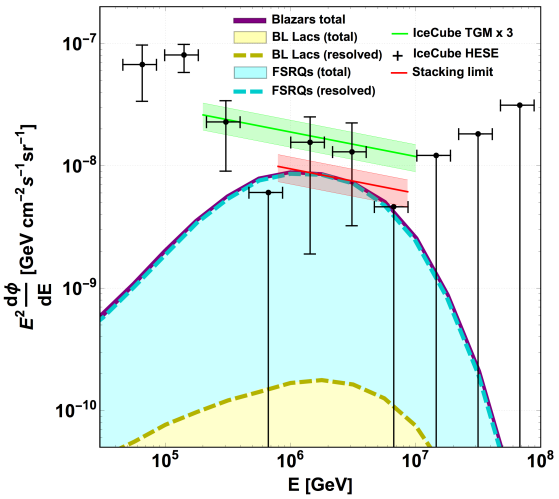

Figure 4.4: Results from scenario 1 (constant baryonic loading ξ_{CR} for all sources. The black points represent the all-flavor fluxes of IceCube high-energy starting events (HESE, Aartsen et al. (2015)), the green line represents the throughgoing muon flux (multiplied by 3 to account for all three neutrino flavors, Aartsen et al. (2016)) and the red line represents the IceCube stacking limit on the blazar contribution (Aartsen et al., 2017a). The shaded yellow and cyan represent the contribution of BL Lacs and FSRQs, respectively, to the total flux of neutrinos, whereas the dotted yellow and cyan curves denote the contribution from resolved sources only. The purple solid curve represents the total neutrino flux expected from blazars. In this case most of the flux is powered by resolved sources, particularly FSRQs. Figure taken from Palladino et al. (2019).

4.2 Results

In this section we analyze the consequences of the hypothesis that the most energetic neutrinos observed by IceCube are produced by blazars. This means that the neutrino flux at the highest energies is saturated by blazars, while at lower energies additional contributions may be present due to the expected peaked nature of the blazar neutrino flux. Although other studies (Aartsen et al., 2017a; Murase et al., 2018; Hooper et al., 2018) disfavor blazars as the dominant source of the astrophysical neutrinos based on different arguments, it is still premature to rule out this source class in light of the potential contribution from unresolved blazars and the fact that they may dominate only at the highest energies. One motivation for such an investigation is that blazars dominate the gamma-ray sky above 100 GeV; if these gamma rays are of hadronic origin, the same sources should efficiently produce neutrinos.

We test three different hypotheses for the baryonic loading and therefore for the relation between the gamma-ray luminosity and the neutrino luminosity defined in Eq. (4.1):

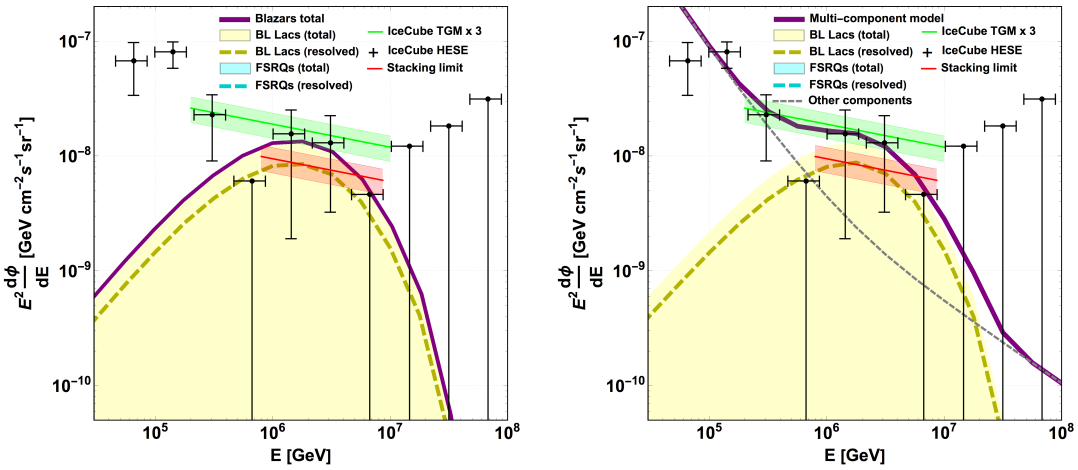

Figure 4.5: Results from scenario 2 (constant ratio L_ν/L_γ). *Left:* theoretical expectation for the neutrino flux, as described in Figure 4.4, compared to the IceCube results. The flux of neutrinos is obtained considering the best fit values of $\epsilon_\nu \times \xi_{\rm CR}$, *i.e.*, $\epsilon_\nu \times \xi_{\rm CR} = 10.5\%$ for low-luminosity blazars, and $\epsilon_\nu \times \xi_{\rm CR} = 0\%$ for high-luminosity blazars. In this scenario, the neutrino flux is powered by both unresolved and resolved sources. *Right:* the same as the middle panel, but including the effect of other possible contributions, such as the residual atmospheric background and Galactic neutrinos, as discussed in Palladino & Winter (2018). Figure taken from Palladino et al. (2019).

- ***Scenario 1: constant baryonic loading*** $\xi_{\rm CR}$ – We assume a constant baryonic loading $\xi_{\rm CR}(L_\gamma) = \bar{\xi}_{\rm CR}$ for all sources, as assumed by Zhang & Li (2017);

- ***Scenario 2: constant ratio*** L_ν/L_γ – We assume that the ratio between neutrino luminosity and gamma-ray luminosity is constant, *i.e.*, $L_\nu/L_\gamma \equiv \mathcal{K} = \text{const}$. This assumption is model-independent and has been used in previous works, such as Wang & Li (2016); Kadler et al. (2016); Righi et al. (2017); Halzen & Kheirandish (2016) to evaluate the flux of neutrinos from BL Lacs. In the context of our model, it implies that the product $\xi_{\rm CR} \times \epsilon_\nu = \mathcal{K}$ (*cf.*, Eq. (4.1)) is constant, which means that $\xi_{\rm CR} \propto (\epsilon_\nu)^{-1}$. We allow for different values for low-luminosity BL Lacs and high-luminosity BL Lacs/FSRQs (for the purposes of this analysis we consider high-luminosity sources as one group), reflecting the potentially different baryonic loadings.

- ***Scenario 3: baryonic loading evolving with*** L_γ – This is a generalization of scenario 2, where we allow the baryonic loading change continuously as a function of L_γ. FOr simplicity, this function is assumed to apply both to BL Lacs and FSRQs.

The theoretical predictions will be compared with the throughgoing muon flux (Aartsen et al., 2016), with the high-energy starting events (Aartsen et al., 2015) above 100 TeV and with the blazar stacking limit (Aartsen et al., 2017a). As discussed in depth in Palladino & Vissani (2016); Palladino et al. (2016); Palladino & Winter (2018), the IceCube data below 100 TeV can be affected by the presence of Galactic neutrinos and residual atmospheric background (both conventional and prompt neutrinos). For this reason, we choose the throughgoing muons as reference for the extragalactic neutrino flux.

In the following three subsections, we discuss the results obtained in the three scenarios described above.

4.2.1 Scenario 1: constant baryonic loading

The diffuse flux obtained choosing a constant baryonic loading is represented in Figure 4.4, where a baryonic loading $\bar{\xi}_{\rm CR} = 150$ has been chosen in order to not overshoot the present IceCube stacking limit for blazars as given by Aartsen et al. (2017a) (specifically, we base ourselves in the stacking limit given by this reference corresponding to a spectral index of $\alpha = 2.2$).

In Figure 4.4, the cumulative flux is represented using a purple solid curve, and the separate contributions from BL Lacs and FSRQs are shown as shaded yellow and cyan regions, respectively. From the same plot we can also see the flux produced by resolved (dashed) and unresolved sources (solid). In this case the total flux is dominated by resolved sources, which is limited by the stacking limit and therefore the observed fluxes cannot be explained. We therefore conclude

that assuming a constant baryonic loading for all blazars, it is not possible to interpret the IceCube observations while respecting the stacking limit. Even if one allowed for two different values for BL Lacs and FSRQs, the result would not change, since the neutrino flux is strongly dominated by the latter.

4.2.2 Scenario 2: constant ratio L_ν/L_γ

As a second possible scenario, we consider a constant value of $\mathcal{K} \equiv L_\nu/L_\gamma$. We assume two different values for low-luminosity ($\mathcal{K}_{\rm low-lum}$) and high-luminosity sources ($\mathcal{K}_{\rm high-lum}$). This *ansatz* is motivated by the fact that if blazars are to power the neutrino flux detected by IceCube without violating the stacking limit,then the contribution of unresolved sources (mainly low-luminosity) has to be enhanced compared to that of high-luminosity sources. As a splitting point, we choose 3.2×10^{45} erg s^{-1} (roughly the separation between the BL Lacs with positive and negative evolution, as shown in the right panel of Figure 2.6). We therefore have:

$$L_\nu = \begin{cases} \mathcal{K}_{\rm low-lum}\, L_\gamma & \text{for} \quad L_\gamma \leq L_\gamma^{\rm evo} \\ \mathcal{K}_{\rm high-lum}\, L_\gamma & \text{for} \quad L_\gamma > L_\gamma^{\rm evo}, \end{cases}$$

with $L_\gamma^{\rm evo} = 3.2 \times 10^{45}$ erg/s.

Since in scenario 2 the value of $\mathcal{K} = \epsilon_\nu \times \xi_{\rm CR}$ is constant for all sources, for our source model this implies that the baryonic loading scales as $\xi_{\rm CR}(L_\gamma) \propto \epsilon_\nu(L_\gamma)^{-1}$,*i.e.*, the baryonic loading $\xi_{\rm CR}$ decreases with luminosity so as to compensate for the increasing neutrino production efficiency $\epsilon_\nu(L_\gamma)$. Note that the scaling of the baryonic loading depends on the particular model of neutrino production in the source, while the assumption $L_\nu \propto L_\gamma$ is independent of the model.

We have computed the total and resolved fluxes from BL Lacs and FSRQs and compare them with the measured throughgoing muon flux (Aartsen et al., 2016) and with the blazar stacking limit (Aartsen et al., 2017a), for a series of pair of values of $\mathcal{K}_{\rm low-lum}$ and $\mathcal{K}_{\rm high-lum}$. We conclude that the best fit is given by the following set of parameters:

$$\mathcal{K}_{\rm low-lum} = 10.5\%; \quad \mathcal{K}_{\rm high-lum} = 0\%. \tag{4.4}$$

This result echoes the previous conclusion that the neutrino flux must be dominated by low-luminosity blazars in order to respect the stacking limit. Within 1σ, a value of $\mathcal{K}_{\rm high-lum} \leq 0.5\%$ is allowed, thus providing an upper limit to the contribution of high-luminosity blazars.

Within scenario 2, we can explain the IceCube data at highest energies without violating the stacking limit, as shown in the left panel of Figure 4.5. From this plot we see that the contribution of FSRQs is absent, while BL Lacs provide the dominant contribution to the diffuse neutrino

flux (yellow shaded region). We then combine this result with the idea presented by Palladino & Winter (2018) that the neutrino flux at lower energies can be explained by other components, such as atmospheric neutrinos below 100 TeV and Galactic neutrinos with an energy cutoff at 150 TeV. The total neutrino flux obtained is shown in the right panel of Figure 4.5. It is important to remark that in principle the atmospheric background should have been already subtracted in the IceCube data points; however, Palladino & Winter (2018) discuss the possibility that a certain residual background can still affect the IceCube measurement. We can see from the right panel of Figure 4.5 that the combination of scenario 2 and these additional components allows for the explanation of the spectrum measured by IceCube.

4.2.3 Scenario 3: baryonic loading evolving with luminosity as a power law

Here we generalize scenario 2, by allowing the baryonic loading to scale with the gamma-ray luminosity as a continuous function, obtained as follows:

- for low-luminosity sources, we maintain the relation $\mathcal{K}_{\text{low-lum}} \equiv L_\nu/L_\gamma = 10.5\%$ favored by the results of scenario 2;
- for high-luminosity sources, we use the upper limit $\mathcal{K}_{\text{high-lum}} \equiv L_\nu/L_\gamma < 0.5\%$ derived in scenario 2;
- we connect these two ranges by making use of the results by Gao et al. (2018), who estimate a baryonic loading of the blazar TXS 0506+056 of value $\xi_{\text{CR}} \simeq 3 \times 10^4$ during the 2017 flare (*cf.* Section 2.5). The inclusion of the result regarding TXS 0506+056 is particularly relevant for this study, since this source has been identified as the first neutrino blazar Section 2.5.

The above information is represented in the left panel of Figure 4.6, where the baryonic loading $\xi_{\text{CR}} = \mathcal{K}/\epsilon_\nu$ is shown as a function of luminosity. The pink band represents the Eddington luminosity (Eq. 2.6) for a range of Lorentz factors $\Gamma_{\text{b}} = 10 - 40$ and a range of black hole masses $M \in [10^{7.5}, 10^{9.5}]M_\odot$, which is typical for blazars (Ghisellini & Tavecchio (2008); Ghisellini et al. (2010), see also Figure 7 of Yu et al. (2015)). The Eddington luminosity is relevant because, as discussed in Section 2.4, we can estimate the maximal plausible baryonic loading by comparing the physical luminosity (Eq. 2.5) of the jet in protons, $L_p^{\text{phys}} = \xi_{\text{CR}} L_\gamma^{\text{phys}}$, with the expected Eddington luminosity of the blazar.

As we can see from this plot, the baryonic loadings of blazars in scenario 3 are compatible with the estimated Eddington luminosities. Note, however, that the Eddington luminosity does not constitute a hard limit on the physical luminosity of blazars, since it may be exceeded during flares (Sadowski & Narayan, 2015).

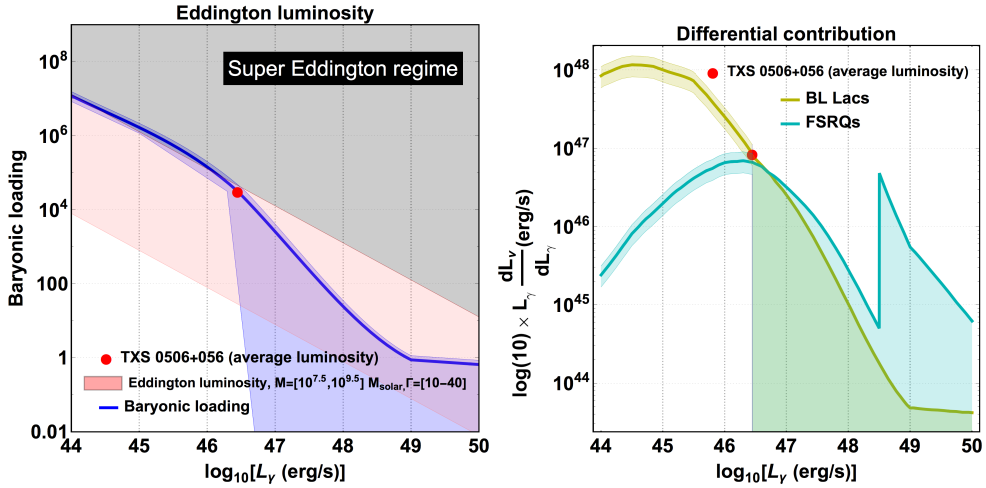

Figure 4.6: Representation of scenario 3. *Left:* baryonic loading as a function of L_γ (the blue curve is the best fit, and the blue band represents the uncertainties). The pink band represents the Eddington luminosity, assuming $\Gamma_b = 10 - 40$ and a black hole mass of $10^{7.5}$-$10^{9.5} M_\odot$. The gray region is disfavored, since the sources would exceed the Eddington luminosity. *Right:* differential contribution to the total neutrino spectrum as a function of gamma-ray luminosity, $L_\gamma dL_\nu / dL_\gamma$, of BL Lacs (yellow) and FSRQs (blue). The discontinuity on the FSRQ curve is due to geometry assumptions of the model (see main text). Figure adapted from Palladino et al. (2019).

The right panel of Figure 4.6 shows the differential contribution $dL_\nu/dL_\gamma$to the neutrino flux from BL Lacs and FSRQs resulting from our baryonic loading function. The differential contribution is obtained as follows:

$$\frac{dL_\nu}{dL_\gamma} = \int_0^{z_{max}} \frac{dz}{(1+z)^2}\; L_\gamma\; \epsilon_\nu(L_\gamma)\; \xi_{\rm CR}(L_\gamma)\; \frac{dN}{dL_\gamma dz}\,, \tag{4.5}$$

where $z_{max} = 6$ both for BL Lacs and FSRQs. From this plot, it is clear that the most important contribution to the high-energy neutrino flux comes from low-luminosity blazars, mostly BL Lacs (yellow-green curve), while the contribution of FSRQs (blue curve) is suppressed like in scenario 2. The discontinuity observed in the FSRQ contribution at about 5×10^{48} erg s^{-1} is due to the geometry assumptions, by which above this luminosity the BLR of FSRQs is large enough that the jet blob is immersed in the BLR, and therefore the external fields can contribute additionally to neutrino production in the jet (see Section 2.4.4, in particular Figure 2.8).

The spectrum resulting from scenario 3 is shown in Figure 4.7. As we can see, like in scenario 2, we can explain the throughgoing muon flux without violating the stacking bound, as we can see. However, this scenario is more flexible insofar as it allows for a small contribution of FSRQs, which was not present in the last scenario. However, as shown by the blue shaded region in the right panel of Figure 4.6, the uncertainties also allow for a null contribution of FSRQs.

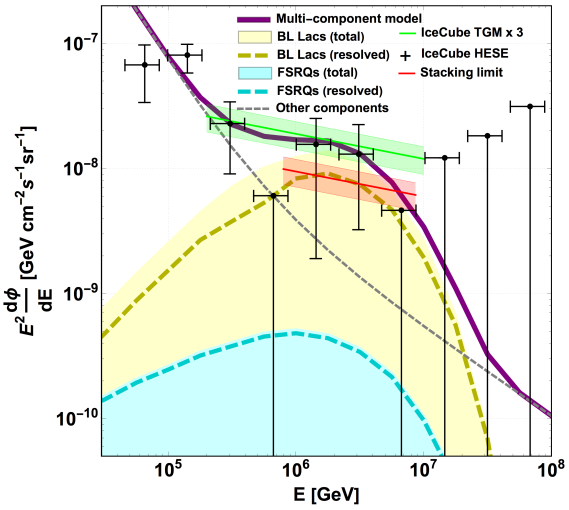

Figure 4.7: Results from scenario 3 (baryonic loading changing continuously with luminosity, see Figure 4.6). As in scenario 2, the neutrino flux is powered by both resolved and unresolved blazars, with low-luminosity BL Lacs providing the dominant contribution. Figure taken from Palladino et al. (2019).

4.3 The case of TXS 0506+056

We here discuss the implications of the neutrino associated with TXS 0506+056 in the context of our model.

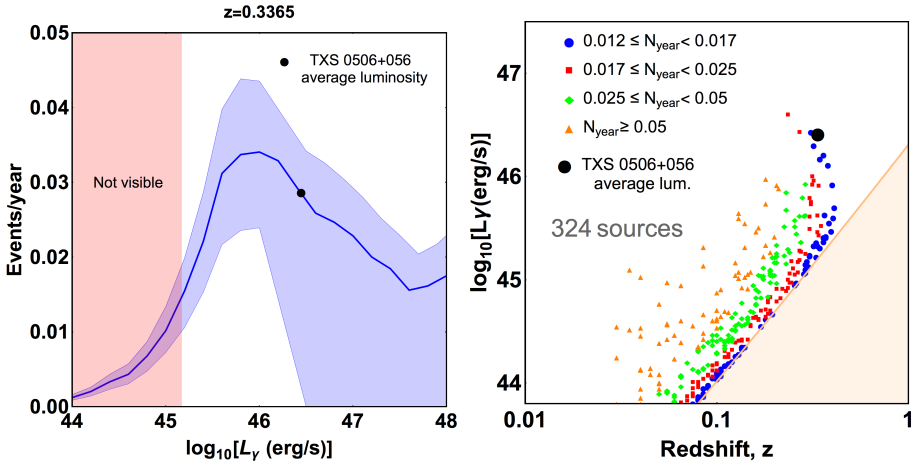

Figure 4.8: *Left:* expected number of events produced by a single source at the same redshift as TXS 0506+056 as a function of the gamma-ray luminosity (event rate refers to point source analysis). The shaded pink region denotes the unresolved sources at redshift $z = 0.3365$. For a TXS-like source, the (time-averaged luminosity) number is equal to 0.012 events per year. *Right:* distribution of resolved sources that can produce at least the same number of events per year as TXS 0506+056. The different symbols indicate ranges of the expected number of events per year from the source, as described in the caption. Figure taken from Palladino et al. (2019).

Expected event rate from TXS 0506+056– According to scenario 3, the luminosity of neutrinos from TXS is expected to be about 1% of the gamma-ray luminosity. This is in agreement with the results by Gao et al. (2018), who modeled the source using a lepto-hadronic hybrid model (*cf.* Section 2.5). Moreover, we can evaluate the neutrino flux expected from TXS following the same procedure used in Section 4.1.1 to evaluate the cumulative neutrino flux expected from the entire blazar population.

Using the effective area for point sources reported by Aartsen et al. (2016), the average rate of events expected from TXS is 0.028 events per year (with deposited energy larger than 100 GeV). In the left panel of Figure 4.8 we show the number of events per year expected from blazars with the same redshift as TXS but different luminosity. From this plot it is clear that blazars with luminosity $10^{45} \lesssim L_\gamma(\mathrm{erg/s}) \lesssim 10^{47}$ are the best candidates to produce high-energy neutrinos. At lower luminosities, the neutrino efficiency $\epsilon(L_\gamma)$ becomes too low (see Figure 4.3), whereas according to our scenario 3, at high luminosities the baryonic loading $\xi_{\mathrm{CR}}(L_\gamma)$ starts to decrease

(Figure 4.7).

In the throughgoing muon dataset, on the other hand, we expect 0.004 events per year from TXS. This estimate is obtained by considering the response function given in Fig. S4-S5 of Aartsen et al. (2018c), and assuming a 2 PeV neutrino energy threshold.

As discussed in Section 2.5, the observation thus far of one event associated to the TXS 0506+056 does not imply that the source must emit 1 event/year. Instead, as argued by Strotjohann et al. (2018), it is possible that there is a large number of faint sources each with a lower individual neutrino output rate, and TXS happened to be the first one to be detected. In fact, the probability of detecting at least one event is given in general by $P(>0) = 1-\exp(-\mu)$, where μ is the expected event rate. Then, for $\mu \ll 1$, the previous expression becomes $P(>0) \simeq \mu$. Within the assumption that there are k sources contributing to the total neutrino flux, the rate of each individual source must be only $\mu \simeq k^{-1}$ in order to yield a non-negligible probability to observe one event from one of the sources.

Expected event rate from the entire blazar population– In the right panel of Figure 4.8, we show the distribution of blazars expected to produce an event rate equal or grater than TXS, using the BL Lac distribution described in Section 2.3.2. We find that the total number of such sources is 324. Using the IceCube point source effective area (Aartsen et al., 2018c) and integrating over luminosity and redshift, we find that each of these sources should produce an average of 0.11 events per year in the point source analysis (above 100 GeV), and 0.016 events per year in the throughgoing muon dataset (above 200 TeV). This is roughly four times higher than that expected from TXS. Therefore, future correlations should be expected at somewhat lower luminosities and redshifts.

In order to estimate the number of events per year from these sources observed by IceCube, the following must be taken into account: assuming that these 324 sources are isotropically distributed, IceCube only detects those that are visible from the Northern hemisphere, via throughgoing muons (*cf.* Section 2.2). Moreover, the alert system is active for neutrinos above 500 TeV (Aartsen et al., 2017c) and at this energy only neutrinos coming from a declination of $0° \leq \delta \leq 30°$ can cross the Earth. Therefore, IceCube is sensitive to roughly one-fourth of the sky (or, equivalently, one-fourth of the sources).

Following these considerations, the expected rate of IceCube events above 100 GeV is roughly given by

$$N_c^{100\mathrm{GeV}} \sim \frac{0.11 \text{ events}}{\text{year} \times \text{source}} \times \frac{324 \text{ sources}}{4} = \frac{8.9 \text{ events}}{\text{year}}. \tag{4.6}$$

These events, however, are not easily distinguishable from the atmospheric background, and it is therefore more interesting to evaluate the rate of events expected in the throughgoing muon

dataset, *i.e.* with energy above 200 TeV. This rate is given by

$$N_c^{200\mathrm{TeV}} \sim \frac{0.016 \text{ events}}{\text{year} \times \text{source}} \times \frac{324 \text{ sources}}{4} = \frac{1.3 \text{ events}}{\text{year}}. \tag{4.7}$$

This rate is consistent with the hypothesis that half of the throughgoing muon flux is produced by resolved blazars, since in 8 years the contribution from those resolved sources would be 10.4 events. Currently, after 8 years of observations, 36 throughgoing muons have been observed (Aartsen et al., 2017d). Given that 2/3 of them are expected to be signal events, we expect roughly 12 events from resolved sources.

Expected number of correlations per year– In order to estimate the rate of correlations per year, *i.e.* the rate of neutrinos associated with flares, we need to refer to the alert system and use the alert effective area. We then obtain an average of 0.004 events/year from each source. Assuming a duty cycle (*i.e.*the fraction of the time that the source is in a flaring state) of 10%, as is the case of TXS 0506+056 (Murase et al., 2018), and also assuming no neutrino production during the quiescent state, then the rate would have to be 10 times higher *during flares*, corresponding to 0.04 events/year. Finally, assuming that only 10% of the sources are in the flaring state at any given time, the expected number of association is given by:

$$N_c^{\mathrm{coinc.}} \sim \frac{0.04 \text{ events}}{\text{year} \times \text{source}} \times \frac{324 \text{ sources}}{4} \times 0.1 = \frac{0.32 \text{ events}}{\text{year}}. \tag{4.8}$$

The alert system has been active for about two years (Aartsen et al., 2017c), and therefore the estimated number of correlations is equal to 0.64. As we know, up to now,only one correlation has been observed which is consistent with our estimate within the Poissonian uncertainty.

4.4 Summary and discussion

In this chapter we studied the possibility that the diffuse astrophysical neutrino flux at sub-PeV to PeV energies is entirely powered by blazars. A major challenge is the stacking limit from the lack of statistical correlations between the arrival direction of neutrinos and the positions of known gamma-ray blazars. This constrains the contribution of resolved blazars to the total neutrino flux, but not that of unresolved sources. Using the blazar sequence and the neutrino production model presented in the previous chapter, we have derived the implications of neutrino observations for blazars.

We demonstrated that the assumption of a constant baryonic loading over the blazar sequence does not allow for a description of the neutrino data, because by fixing the baryonic loading, high-luminosity objects will dominate the neutrino flux, and therefore the stacking limit will

constrain the blazar contribution to the diffuse flux. In order to circumvent this limitation, we have allowed the baryonic loading to change as a function of luminosity. We analyzed two different possibilities: in the first one the ratio between luminosity in neutrinos and luminosity in gamma rays is constant, and in the second one the baryonic loading scales continuously with the gamma-ray luminosity.

We have found that the only scenario in which blazars can explain the high-energy neutrino flux is that where the baryonic loading of low-luminosity objects is higher than 10^5, whereas the baryonic loading of high-luminosity sources (both BL Lacs and FSRQs) must be lower than ~ 100. Under this hypothesis, low-luminosity objects can power the entire neutrino flux above a few hundreds of TeV, while the contribution of high-luminosity objects, particularly FSRQs, is highly suppressed. We demonstrated that in this scenario the baryonic loading of blazars of all luminosities satisfies the Eddington limit. While previous works have indicated that such high baryonic loadings are difficult to be achieved in low-luminosity BL Lacs because of their relatively inefficient accretion, we may infer that the conditions are different if the neutrino emission occurs during flares with baryon injection.

The recent observation of neutrinos from TXS 0506+056 can also be interpreted in our model. We find that the average rate of neutrino events from an object with the same luminosity and redshift as TXS is of about 0.004 throughgoing muons per year. Taking into account the blazar luminosity function, we find on average 0.016 events per blazar per year, coming from about 300 objects resolved by *Fermi* that are at least as good neutrino emitters as TXS 0506+056. This yields about 0.3 expected associations per year between neutrinos and known blazars. These associations are likely to come from BL Lacs with luminosities $L_\gamma \sim 10^{45}$ erg s^{-1} and redshifts $z \sim 0.1$.

In conclusion, we have demonstrated that the observed flux of throughgoing muons is well-described, without violating the stacking limit for blazars, if one accepts large enough baryonic loadings for unresolved (low-luminosity) sources, while high-luminosity BL Lacs and FSRQs are disfavored as the dominant contributors to the diffuse flux of high-energy neutrinos.

Chapter 5

Modeling the 2014–15 neutrino flare of TXS 0506+056

As discussed in Section 2.5, in September 2017 a muon neutrino with energy 290 TeV was observed by IceCube from a position compatible with the blazar TXS 0506+056 (referred to throughout this document as TXS). The neutrino was detected in temporal coincidence with a flare in multiple wavelengths (Aartsen et al., 2018b). This event has enticed the multi-messenger community to explore the potential of TXS as a source of astrophysical neutrinos. In this chapter, we study the possibility that the neutrino excess observed from the same direction in a six-month period in 2014–15 (Aartsen et al., 2018c, henceforth referred to as the *historical* neutrino flare) can have originated in photo-hadronic interactions in the same source. Compared to the 2017 flare, explaining the 2014–15 signal implies accommodating a high enough neutrino flux with the lack of electromagnetic activity (Murase et al., 2018, *cf.* Section 2.5). For that purpose, we draw from some ideas applied by other authors to explain the 2017 flare within a photo-hadronic scenario (especially the models by Gao et al., 2018; Cerruti et al., 2018b; Keivani et al., 2018, see Section 2.5 for a review of other geometries). In particular, we test a one-zone model, a compact core model, and an external field model. The latter is particularly relevant since TXS has recently been argued to possess FSRQ-like external emission (Padovani et al., 2019).

The results discussed in this chapter are based on the work by Rodrigues et al. (2018b). My contribution was the development of the external field model, using the numerical code AM^3 (*Astrophysical Modeling with Multiple Messengers*), applied previously by Gao et al. (2018) to the 2017 TXS flare. The details about the AM^3 code can be found in the original publication by Gao et al. (2017). For this work, I extended the AM^3 code to include interactions with external fields from BLR emission, treated as the injection of a static photon spectrum into the blob. Furthermore, I created a genetic algorithm to perform efficient scans of the physical parameters in order to find the best result for the external field model, *i.e.* a parameter set that can account for the maximum number of IceCube events while best obeying the constraints from multi-wavelength observations.

5.1 Neutrino observations and multi-wavelength constraints

Neutrino signal– Triggered by the multi-messenger discovery of the 2017 flare, IceCube searched their archival data for an excess from the direction of TXS (Aartsen et al., 2018c) for the entire duration of IceCube's data taking. In the period between October 2014 and March 2015, a temporal clustering has been detected of 64 events in total within 3° of the direction of TXS. By using a likelihood function, in which the atmospheric background is taken into account and the signal is assumed to be distributed as a power law, a 3.5 σ excess over the atmospheric background was found by Aartsen et al. (2018c), with an estimated number of signal events of 13 ± 5. The most energetic event has a deposited energy in the detector of 20 TeV, while most events have energies around $\sim$ 10 TeV.

One of the most important results of the models that will be tested is the expected number of muon tracks in IceCube. That number is computed by convolving the neutrino spectrum predicted by the model with the effective area of IceCube Aartsen et al. (2018c) (IC86b data period) at the given neutrino energy and for the declination of the source (about 5.7°). If the model parameters appropriately describe the neutrino signal, this number should be close to the observed number of 13 ± 5 events. Because the signal consists only of muon tracks (which have good angular resolution, *cf.* Section 2.2), only muon neutrinos are of interest. If the neutrinos are produced by photo-meson production, as assumed here, then given Eq. (2.11) they are emitted from the source as both muon and electron neutrinos, with a number ratio of 2:1. For this particular ratio, during propagation to Earth the neutrinos become equally distributed among all three flavors (electron, muon and tau), due to oscillations (see *e.g.* Becker, 2008). The flux of muon neutrinos observed is therefore approximately one-third of the total neutrino flux emitted from the source.

Multi-wavelength constraints– Interestingly, the neutrino signal was not accompanied by any significant increase in electromagnetic emission. In contrast to the 2017 flare, the multi-wavelength data from this period are very sparse and the only constraints on the SED can be derived from gamma-ray flux measurements by *Fermi*-LAT (Garrappa, 2018), as well as the radio and optical monitoring data compiled by Padovani et al. (2018). Additionally, the *Swift* Burst Alert Telescope (BAT) that monitors X-ray transients and performs regular sky surveys, was not triggered during the period of the neutrino flare and did not detect TXS in the 15–50 keV band, implying that its flux during the neutrino flare was significantly less than 3 mCrab, or 7.2×10^{-11}erg cm^{-2} s^{-1} (Krimm et al., 2013). Based on *Fermi* data, Padovani et al. (2018) have speculated that there may be a hardening in the SED of the source above 2 GeV during the neutrino flare, although this feature may in fact not be significant (Garrappa, 2018).

In this study we use the available observational evidence during the neutrino emission period

in order to constrain the source model. Given the limited constraints from long-term photon data, we do not attempt to derive an explanation for the time-dependent behavior, namely the transition between the neutrino-bright and neutrino-quiet states.

5.2 One-zone and compact core models

The models for the neutrino and photon emission are constructed using the lepto-hadronic code AM^3, which solves self-consistently the time-dependent kinetic equations for non-thermal protons, neutrons, electrons, positrons, photons and neutrinos produced in the jet. Unlike NeuCosmA, used in Chapters 3 and 4, here the photons are computed self-consistently as products of the leptonic and hadronic interactions. This feature is essential in order to estimate the emitted photon fluxes and evaluate the model against observations. All the interactions introduced in Section 2.3.1 are included self-consistently in this calculation. An example of a result obtained using this code was already shown previously in Figures 2.7 and 2.10.

For the one-zone and compact core models, the parameter scans were performed using a grid search. On the contrary, as discussed in the next section, the results of the external field scenario were obtained using a more sophisticated optimization algorithm. Note that in spite of the extensiveness of the parameter space scan, we still cannot claim completeness of the results due to the complexity of the problem.

Firstly, a one-zone model was tested (Section 2.4.3). Unlike the treatment in Chapter 3, here we consider the particle escape timescale to be independent of energy, given in terms of the light-crossing time of the blob, $R'_{\rm blob}/c$, as $t'_{\rm esc} = R'_{\rm blob}/(c\, f_{\rm esc})$. In the one-zone model (as well as the compact core model), the constant $f_{\rm esc}$ was fixed to 3×10^{-3}. The parameters that were scanned are the size of the blob $R'_{\rm blob}$ and its bulk Lorentz factor and Lorentz factor $\Gamma_{\rm b}$, the magnetic field strength B', the maximum Lorentz factors of the injected electrons and protons, $\gamma_e^{\rm max}$, $\gamma_p^{\rm max}$, and their respective total injection luminosities, L'_e and L'_p. On the contrary, for simplicity, the spectral index of the proton and electron injection spectra has been fixed to 2, as well as in the other models in this study.

Upon scanning these parameters, we obtained two distinct "classes" of parameter sets, whose SEDs and neutrino spectra are represented in Figure 5.1: in the left plot, we show a set of solutions where the emitted SED is in reasonable agreement with the multi-wavelength constraints; however, the number of emitted neutrinos expected during the duration of the flare is at most 1.8 (in the case depicted in red, whose parameters are listed in Tab. A.2). This low number of neutrinos compared to the IceCube signal can be seen by comparing the predicted spectra (dashed curves) with the expected flux given by Aartsen et al. (2018b) (blue line and shaded area). Note that the shape of the signal assumes a power law neutrino spectrum, which as

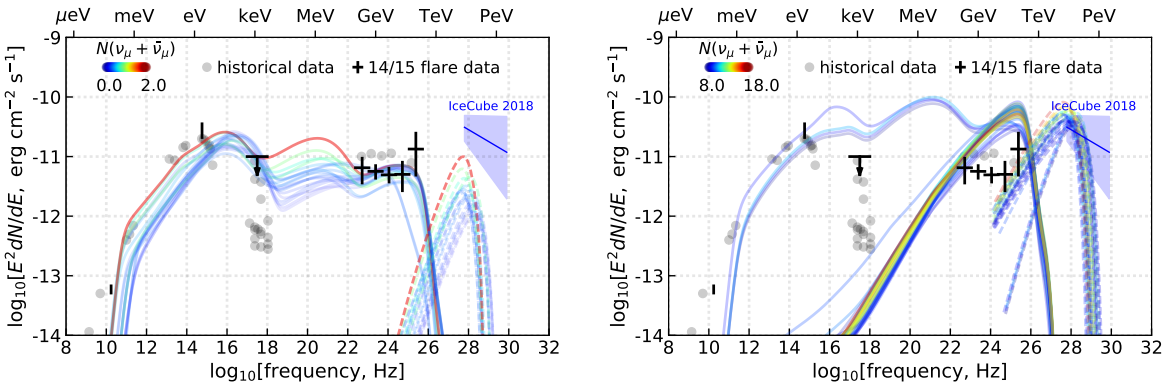

Figure 5.1: Spectral energy distributions and muon neutrino fluxes predicted by the one-zone hadronic model, compared to the single-flavor flux derived by IceCube (Aartsen et al., 2018c). In the left panel, the parameter sets are optimized to describe the SED in agreement with observations, while failing to explain neutrino emission; in the right panel the parameter sets account for 13 ± 5 of the observed muon neutrinos, but overshoot the multi-wavelength emission such as in the *Fermi*-LAT range. In Tab. A.2 we show the list the parameters for the red curve in the left panel, and the ranges of parameters for all the spectra in the right panel. The data used to test the SED are also plotted with the respective error bars in black (see main text), where only data during the neutrino flare are used, and historical data taken during the years before 2017 from the database of the Space Science Data Center and from the NASA/IPAC Extragalactic Database are shown in gray for reference. The radio data were also recorded simultaneously with the neutrino flare (Padovani et al., 2018). Figure taken from Rodrigues et al. (2018b).

we know is not the case for photo-hadronic sources; we therefore do not expect our predicted neutrino spectra models to fit this spectral shape, but focus only on the number of neutrinos. Regarding the SEDs, the first bump come from synchrotron and IC emission off $e^{\pm}$ that originate in $\gamma\gamma$ annihilation, and the MeV bump from Bethe-Heitler pair production (compare this to the case shown in Figure 2.7). This example demonstrates the importance of electromagnetic data across the entire spectrum to constrain theoretical models, since the electromagnetic cascade accompanying the neutrino production can be hidden in unconstrained energy bands like the MeV range.

On the other hand, in the context of a one-zone model, demanding that the number of IceCube events be within 1 σ of the observed signal leads to an SED that is in tension with observations, as seen in the right panel of Figure 5.1. In Tab. A.2 we list the parameters of the solution shown in yellow this plot, which yields 13.2 neutrino events. In this case, we see that the efficient photo-meson production leads to the development of strong hadronic cascades in the source that overshoot the X-ray constraint and the gamma-ray observations. We also find a cluster of SEDs that are strongly dominated by IC scattering, due to a compact emission region and weak magnetic field. This leads to hard gamma-ray emission, high $p\gamma$ efficiency and higher neutrino fluxes, but the emission at larger wavelengths is highly suppressed, which fails to explain radio and optical observations. On the other hand, the subset of models with sufficient synchrotron emission all overshoot the X-ray and gamma-ray constraints.

Given these results, we conclude that one-zone models of the historical flare are in tension with observations.

As discussed in Section 2.4.5, a smaller emission region within the blob can enhance neutrino production. Drawing from the model used by Gao et al. (2018) to describe the 2017 TXS flare (*cf.* Section 2.5), we speculate that during the 2014–15 neutrino flare a compact core was formed in the jet in addition to the larger blob, and we assume the same Lorentz factor for both emission regions (unlike spine-sheath models for example). A further simplification is that we do not consider the interaction between the non-thermal emission of the blob and the core; the total emission of photons and neutrinos is assumed to originate from the superposition of both radiation zones.

Due to its small size, the core is highly efficient in both IC scattering and photo-meson production, because of the high density of electrons and of target photons for either of these interactions. On the other hand, synchrotron emission depends on the total number of electrons (Eq. 2.8) and not on the compactness of the source, which implies that synchrotron emission from the core can be suppressed; we therefore rely on the larger blob zone for the emission of radio to optical.

The SED and neutrinos for an optimized parameter set of the compact core model are shown

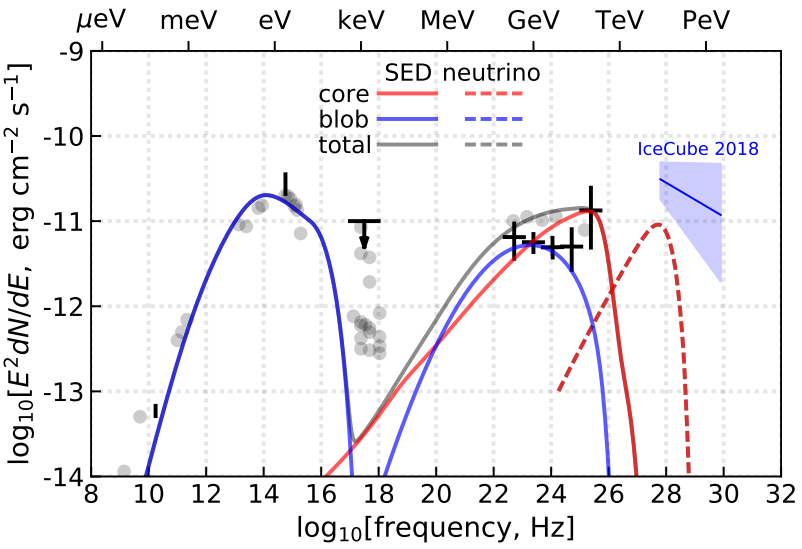

Figure 5.2: Emitted SED and muon neutrino spectrum from TXS for one parameter set of the compact core model (*cf.* Tab. A.2). We plot the contributions from the blob region (blue), which accounts for the fluxes from radio to X-rays but does not produce neutrinos because it is of dominantly leptonic origin, and the contribution from the core region (red), which accounts for the neutrino flux and the gamma-ray data, separately. The parameter set was obtained through optimization of the emitted neutrino flux, which yields at most 1.9 IceCube events. Figure taken from Rodrigues et al. (2018b).

in Figure 5.2, and the respective parameters are appended in Tab. A.2. The neutrino flux emitted by the core translates to 1.9 observed muon tracks, which is slightly higher than can be achieved in one-zone models within an acceptable SED; however, it is still in tension with the IceCube result. Due to the reasons discussed previously, the blob emission describes the data from radio to soft gamma rays, while its contribution above GeV is low. The large volume of the blob and correspondingly low target photon densities for photo-hadronic interactions also make it an inefficient neutrino emitter; the interactions in the blob basically correspond to a purely SSC scenario. On the other hand, the efficient IC scattering taking place in the core leads to bright gamma-ray emission, that dominates over the blob emission above 10 GeV. Note that X-ray emission is highly suppressed, which makes this model viable even if there were more constraining X-ray limits. Solutions can also be found yielding a higher number of neutrino events, but the parameters in question are in general unphysical: for instance, the magnetic field in the blob could in principle be stronger than that in the core, but this would be physically unrealistic, since the core is the more compact region.

5.3 External field model

We now focus on the external thermal field model, which I developed based on the same AM^3 code used for the one-zone and compact core models discussed previously. Although TXS is observationally a BL Lac due to the absence of atomic or molecular broad lines (*cf.* Section 2.3.1), Padovani et al. (2019) have recently noted that there is in fact some indirect observational evidence that may support the existence of external fields like those of FSRQs. This justifies an external field model, like that applied by Keivani et al. (2018) to the 2017 TXS flare.

As explained in detail in Section 2.4.4, a fraction of the disk radiation may be isotropized through Thomson scattering in a BLR surrounding the disk and subsequently Lorentz-boosted into the blob, it the latter is located within the BLR. For simplicity, the thermal continuum is considered here as a simple thermal distribution instead of the multi-temperature spectrum in Section 3.5 (green curves in Figure 2.9). Moreover, the atomic emission is modeled as a single broad line, the hydrogen Ly α line, which as discussed previously is typically the strongest (Greene & Ho, 2005, *cf.* black peaks in Figure 2.9). This leads in general to stronger absorption of VHE gamma rays due to photon annihilation, while boosting neutrino production (see Figure 2.10 and discussion thereafter). After escaping the jet, gamma rays will be further absorbed during their propagation through the BLR (*cf.* orange shaded area in Figure 2.10). Furthermore, given the source's redshift, we know that photons with energy $E > 300$ GeV will suffer additional attenuation due to the interaction with the EBL (purple shaded area in Figure 2.10).

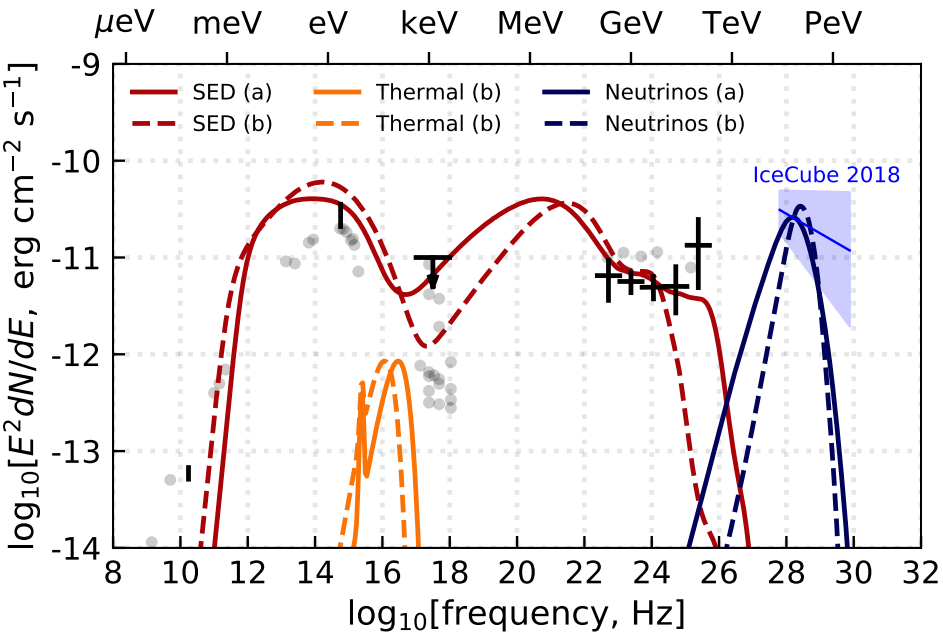

Figure 5.3: Emitted SED (red) and muon neutrino spectrum (blue) from TXS when photon fields from an accretion disk and broad line emission (orange curves, shown in the observer's frame) are included in the calculation. The parameter sets (a) and (b), listed in Tab. A.2, are shown as solid and dashed curves, respectively; the corresponding number of IceCube muon tracks is 4.9 and 4.0. Figure taken from Rodrigues et al. (2018b).

Physical parameters– The external field model has more parameters than the former two, since it includes thermal and broad line emission; however, some of these parameters are constrained by observations. The disk luminosity of TXS has been argued to be around $L_{\rm disk} \sim 5 \times 10^{44}$ erg s^{-1} by Padovani et al. (2019), given a black hole mass of 3×10^8 $M_\odot$. From there, the broad line luminosity can be deduced from empirical relationships (Kaspi et al., 2005) (as discussed in Section 2.4.4, this value is about 10% of the disk luminosity, $L_{\rm disk} \sim 5 \times 10^{43}$ erg s^{-1}). Furthermore, the temperature of the disk of most quasars is observed to lie around a few times 10^5 K (Bonning et al., 2007), similar to the value assumed in Chapter 3 for the peak of the thermal continuum. Because the objective is to maximize neutrino production, the maximum proton Lorentz factor $\gamma_p'^{\rm max}$ has also been fixed so that the maximum-energy protons interact through photo-pion production with the external thermal field at threshold. The maximum electron energy has been fixed so that their characteristic synchrotron frequency coincides with that of the optical data point (however, as we will see, this feature is not relevant since the optical emission is dominated by secondary electrons from hadronic interactions). This leaves *seven* free parameters to be scanned: B', $R'_{\rm blob}$, $\Gamma_{\rm b}$, L'_p, $\gamma_p'^{\rm max}$, $f_{\rm esc}$, and $R_{\rm BLR}$. Although the last parameter is also constrained by observations (Kaspi, 2001; Padovani et al., 2019), there are relatively large statistical uncertainties in these measurements, and allowing for a variation of the BLR radius gives the model an additional degree of flexibility, insofar as the density of external photons in the blob frame depends on this quantity (Eq. 2.18).

Parameter space search method– The scan of this parameter space was performed using an optimization code that I developed for this project, based on a genetic algorithm (see *e.g.* Goldberg, 1989). The most relevant aspects of the method are detailed in Appendix C. In summary, a pool of ten thousand parameter sets are randomly created, and the emitted SED and neutrino spectrum for each parameter set is computed independently and in parallel using the AM^3 code. Each parameter set is evaluated based on the SED and neutrino spectrum. The SED is evaluated through a χ^2 test against the optical and gamma-ray data points and the X-ray upper limit (if the X-ray flux overshoots the observational limit it contributes to the total χ^2 value, while undershooting the limit is not penalized). The number of neutrino events in IceCube obtained from the neutrino spectrum is also evaluated in a similar fashion against an expected signal of 13 ± 5 events (see Appendix C for details). Upon evaluating the entire population of parameter sets, they are allowed to compete for remaining in the next generation, based on their "grade". Then, random mutations are applied in order to probe more effectively the entire parameter space, as well as random crossovers, *i.e.* interchange of one parameter between random pairs of sets. The new parameter sets are evaluated again by running the photo-hadronic code, and the process repeats. The "population" of all parameter sets eventually converges to small areas of the parameter space that yield a high value of the quality criteria.

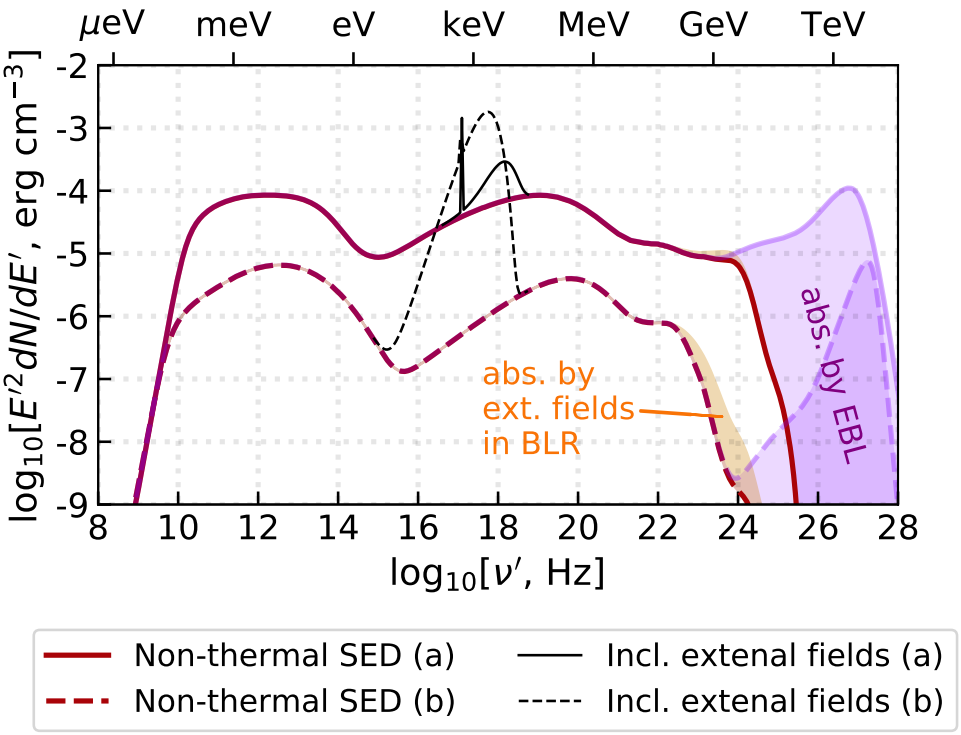

Figure 5.4: Photon energy density spectra in the rest frame of the jet for the two examples shown in Figure 5.3. The external fields, whose contribution is shown in black, are not seen in the observer's frame because they are not isotopic in the jet frame. We also show the effect of gamma-ray attenuation during propagation in the BLR (orange shaded area) and during propagation in the intergalactic medium (purple shaded area).

For this particular problem, the algorithm was repeated for 50 generations, and at the end the best parameter set was taken as the optimized result. The specific methods used for each part of this algorithm are detailed in Appendix C.

Results– In Figure 5.3 we present two optimized results of the external field model, both obtained with the genetic algorithm in separate searches: example *(a)*, shown as solid curves, was obtained by fixing the disk temperature to 5×10^5 K, and example *(b)* by fixing it to 2×10^5 K (dashed curves). Both values are within the statistical uncertainty of quasar observations (Bonning et al., 2007). Note that the emission from the disk and broad line (orange) is out-shined by the highly beamed jet radiation (red) and is therefore invisible to the observer, who sees the object as a BL Lac, as proposed by Padovani et al. (2019). While the entire luminosity

of the hydrogen broad line is beamed into the jet rest frame, the thermal bump that is shown corresponds to the entire disk emission, of which only 1% is considered to be back-scattered into the jet. Both parameter sets are detailed in Tab. A.2; *(a)* leads to 4.9 neutrino events during the flare, and *(b)* to 4.0. As mentioned earlier, in this model the optical emission is dominated by synchrotron emission of secondary $e^{\pm}$ pairs from Bethe-Heitler production, while the bump at a few MeV is emitted by $e^{\pm}$ pairs from annihilation of VHE hadronic photons. Emission from primary electrons is therefore sub-dominant across the entire spectrum.

Apart from the disk temperature, the key differences between example *(a)* and *(b)* concerning hadronic interactions may in fact not be clear from the SED and neutrino spectra. In the left panel of Figure 5.4 we show both SEDs, this time in the jet frame and given as a density spectrum, rather than luminosity (the difference being a constant factor of $4\pi R'^2 c$, see Eq. 2.14). This is why the non-thermal photons (red curves) have very different densities in the jet frame in both examples, while the emitted photon fluxes (Figure 5.3) are at a similar level. We plot in Figure 5.4 the additional contribution of the external field in the jet frame (thin black curves), which does not appear in the observed SED due to the fact that it is not isotropic in the jet frame. As we can see, in the two examples the external photons have very different energy densities, although the disk luminosity is the same in both cases; this is due to the different values of $R_{\rm BLR}$ (*cf.* Eq. 2.18).

As mentioned previously, in both examples the proton maximum energy was determined so that the protons interact at threshold with the external field, resulting in neutrinos with energy in the observer frame[1] of $E_\nu \sim 100\ {\rm TeV}(T_{\rm disk}/10^5\ {\rm K})^{-1}$, and emission of gamma-rays with energy $E_\gamma \sim 200$ TeV (*cf.* discussion below Eq. 2.11). These photons resulting directly from $p\gamma$ interactions are shown as a purple band in Figure 5.4, since they are attenuated by the EBL during propagation. We can see that these multi-TeV photon spectra have very different shapes; namely, the one from example *(a)* is broader, while that from example *(b)* is more peaked around the maximum emission frequency. This is also the case of the emitted neutrino spectra (Figure 5.3). This is because in example *(b)* the target photons for $p\gamma$ interactions are almost entirely provided by the external emission (compare the relative height of the red and black dashed curves of Figure 5.4), which is highly peaked. On the contrary, in example *(a)* the non-thermal radiation has an energy density that is comparable with that from the external fields, due to strong cascades from photo-hadronic interactions. The flatter shape of the non-thermal spectrum then leads to a broader neutrino spectrum. This is consistent with the slower escape of protons and electrons assumed in example *(a)* compared to *(b)* (see Tab. A.2), which leads to

[1]The energy in the observer's frame of the neutrinos produced by $p\gamma$ interactions with the external fields, $E_\nu = E'_\nu \Gamma_{\rm b}/(1+z)$, is in fact independent from the bulk Lorentz factor of the jet, since the target photons are themselves Lorentz-boosted from the black hole frame into the jet frame, $E'_\nu \propto E'^{-1}_{\gamma\ {\rm ext}} = (E_{\gamma\ {\rm ext}}\Gamma_{\rm b})^{-1}$.

a more efficient cascade process due to the continuous presence in the source of the secondary electrons. On the other hand, the faster escape rate of protons in example *(b)* leads to the requirement of a higher proton injection power, which may exceed the Eddington luminosity of the source depending on the assumed black hole mass (see Section 5.4).

The high neutrino production efficiency of this model implies that the gamma-ray spectrum above GeV energies suffers a softening in example *(a)* and a strong cutoff in example *(b)*, due to photon annihilation (at an energy of $E_\gamma \sim 10\ \mathrm{GeV}(T_{\mathrm{disk}}/10^5\ \mathrm{K})^{-1}$ in the observer's frame, *cf.* Figure 5.4). One can therefore say that neutrino efficiency and gamma ray emission are inversely correlated in this model; this has also been concluded in other recent works (Murase et al., 2016; Reimer et al., 2018). This creates tension with *Fermi* observations (Garrappa, 2018; Padovani et al., 2018), as can be seen in Figure 5.3. In spite of this shortcoming of the model, result *(a)* only fails to describe one *Fermi* data point (at ~ 100 GeV), while it is in agreement with observations at other wavelengths within a 1 σ uncertainty; furthermore, note that the *Fermi* data point at ~ 100 GeV consists only of two photons (Garrappa, 2018), and therefore has a large statistical uncertainty. On the other hand, the X-ray bound is less saturated in example *(b)*, since the in-source cascades are suppressed by a faster proton and electron escape rate, as discussed above.

5.4 Summary and discussion

This chapter explored the viability of lepto-hadronic models to explain the 2014–15 neutrino flare from the blazar TXS 0506+056, following Rodrigues et al. (2018b). The space of physical parameters of the source was scanned to find which parameter values, if any, can best explain the observed neutrino signal while obeying the constraints of the sparse observations in optical, X-rays and gamma rays during the neutrino flare. Three distinct assumptions about the geometry of the blazar were tested: a one-zone model, a compact core model similar to that considered by Gao et al. (2018) to explain the 2017 TXS flare, and an external field model, in similar lines to the model by Keivani et al. (2018) to explain the 2017 flare. My main contribution to this study was the development of the external field model, and the optimization of the physical parameters of the model through an automated scanning method.

While none of the models predicts a number of muon track events in IceCube compatible with the observed signal of 13 ± 5 events (Aartsen et al., 2018c), the external field model can explain a maximum of 5 events without conflicting with the optical and X-ray constraints. The results of two optimized parameter sets were shown (Figure 5.3), representative of two qualitatively different scenarios: in one of them (dashed curves), the external photons serve directly as the target field for the interactions, leading to the emission of a highly peaked neutrino spectrum, low

X-ray emission from suppressed cascades and strong gamma-ray attenuation from the external field, leading to a sharp cutoff at 10 GeV. In the other example (solid curves), the external fields leads to strong cascades in the jet, increasing substantially the keV to MeV fluxes (while respecting the X-ray constraint from the *Swift* BAT), and leading to an overall flatter SED and a broader neutrino spectrum. This last model predicts a number of 5 muon neutrinos observed in IceCube, and is therefore in tension with the IceCube signal at a level of 1.6 σ. On the other hand, the excess of 13 ± 5 has been obtained by IceCube under the assumption of a power law neutrino spectrum, and it is still unclear whether the signal would remain the same if the analysis were performed considering instead a more peaked spectrum obtained with a self-consistent target photon field, such as that shown in Figure 5.3.

While an external field model is supported by observations of the source that seem to imply the presence of external fields like in the case of FSRQs (Padovani et al., 2019), these fields may in fact pose a challenge in explaining the variability of the source in gamma rays (such as during the 2017 flare), since continuous high-energy photon attenuation would be expected. On the other hand, this might potentially be indicative of radiation zones in the jet that form at different distances to the black hole: when lying outside the BLR, the external fields would not play a role in non-thermal interactions and the source could be transparent to gamma rays.

The other models that were tested, namely a one-zone and a compact core model, yield lower neutrino fluxes. Nonetheless, the compact core model presents itself as an interesting scenario due to the low X-ray fluxes that are predicted (more than two orders of magnitude below the level of GeV gamma-rays), and the possibility of emission of a hard gamma-ray spectrum. Both of these features are a consequence of the decoupling of the low- and high-energy emission, which are assumed to be produced in distinct zones which do not interact with one another: gamma rays and neutrinos both originate from the core (where IC scattering and $p\gamma$ interactions are very efficient due to the high radiation and electron density), and the optical emission originates in a larger blob via electron synchrotron. The decoupling of the two zones can also account naturally for a fast gamma-ray variability compared to emission in larger wavelengths, which is a general feature of two-zone models. However, in order to explain a multi-wavelength flare like that of 2017, the compact core model would require a coordinated evolution of the parameters of the blob and the core to explain the simultaneous rise and decay of the optical, X-ray and gamma-ray fluxes. Another limitation is that it may be difficult to justify the non-interaction of radiation from the core with that from the blob if the two zones are contained one in the other. A natural solution would be to assume that these regions lie in different points along the jet; however, a large separation should imply temporal delays between the respective emission bands that have not been observed.

Note that unlike the single neutrino observed during the 2017 flare, which is subject to the

Eddington bias (Strotjohann et al., 2018, see Section 2.5), the same argument does not apply to the 2014-15 flare, where the predicted event number needs to be significantly larger than unity in order to be compatible with observations. Since for blazars the number density of target photons (X-rays in the case of TXS) is relatively low compared to other objects such as GRBs, the optical depth to $p\gamma$ interactions is typically much lower than unity – which needs to be compensated for with a large baryonic loading in order to become a significant neutrino source. The photo-hadronic interaction rate can then be enhanced by assuming a smaller production region (like in the compact core model), or a higher density in X-rays (like in the external radiation field model).

This chapter demonstrates the challenges involved in the simultaneous description of the SED and neutrino emission from TXS during the period of the 2014-15 neutrino flare. Besides the difficulty in explaining the high neutrino flux necessary to explain the signal, another major challenge to photo-hadronic models is to describe the different states of the same source, for instance the transition between neutrino-quiet and neutrino-loud states, or the explanation of two different events like the 2014-15 and 2017 flares. However, from our model and the extensive parameter scans performed, we can conclude that *a)* obtaining more than 5 neutrino events during the flare would imply violating multi-wavelength constraints, especially in gamma and X-rays, and that *b)* in the photo-hadronic paradigm, a transition between neutrino-quiescent and flaring states should imply distinctive electromagnetic activity.

Chapter 6

Neutron star merger remnants as sources of cosmic rays

As mentioned in the introduction, the first binary neutrino star merger event was detected August 17, 2017 in gravitational waves (Abbott et al., 2017a), in an event that was named GW170817. Following this detection, a short GRB was detected both by the *Fermi* GBM and the INTEGRAL-SPI detector (Abbott et al., 2017b; Savchenko et al., 2017). Subsequently, an optical counterpart was also detected, dubbed event EM170817 (Arcavi et al., 2017; Coulter et al., 2017; Lipunov et al., 2017; Soares-Santos et al., 2017; Valenti et al., 2017), associated to a kilonova from the GRB remnant. The distance to the remnant estimated from the GW signal, D ∼40 Mpc, was found to be consistent with that of the nearby galaxy NGC 4993 (Hjorth et al., 2017), leading to the conclusion that this was the host galaxy of the merger event that produced the GRB. The subsequent observation of non-thermal emission from the remnant in both radio (Hallinan et al., 2017; Alexander et al., 2017) and X-ray bands (Margutti et al., 2017; Ruan et al., 2018) indicates that particle acceleration is occurring efficiently in the source. This emission was observed to brighten as a power-law in time for a timescale of ∼160 days, with the apparent recent onset of a dimming of the source (D'Avanzo et al., 2018; Nynka et al., 2018).

We will start by constraining physical properties of the remnant of GW170817, such as the magnetic field strength and the energy spectrum of the electron population, based on the observed non-thermal emission from radio to X-rays, assuming it is emitted by electron synchrotron. We then model the gamma-ray spectrum expected from inverse Compton scattering, by means of a semi-analytic leptonic radiation model. We show that the non-observation of the remnant in gamma rays can place a lower limit on the magnetic field strength in the source. Finally, we estimate the maximum energy of CRs accelerated in the source, and its optical thickness to photo-hadronic processes like photo-disintegration. Within the parameter ranges obtained previously, we show that this source class can accelerate CRs to energies between PeV and 10 EeV, *i.e.* between the knee and ankle of the CR spectrum (*cf.* Figure 2.1).

This chapter is based on Rodrigues et al. (2019), to which I contributed with the leptonic modeling of the source and the analysis of the model's results.

6.1 Source energetics

The total kinetic energy released in the outflow can be estimated from the gravitational binding energy of the neutron star binary, and should be of the order of 10^{50-52} erg. Approximating the total energy as 10^{51} erg, and taking also into consideration a total mass of ejected material of $M = 10^{-2}\ M_\odot$ (Takami et al., 2014), these values are compatible with an ejecta velocity of $\beta = V/c = 0.2$. Moreover, measurements of the remnant's photospheric velocity also support non-relativistic speeds ($\beta \lesssim 0.2$ Piro & Kollmeier, 2018)[1]. This mass corresponds to a total number of electrons in the outflow of $N_e = 10^{55}$.

As this material expands, the source eventually transitions from being optically thick to optically thin to its own radiation. The optical thickness to Thomson scattering is given by

$$\tau_{e\gamma} \approx \frac{N_e\,\sigma_T}{4(\beta ct)^2} \approx \left(\frac{t}{2\ \text{days}}\right)^{-2}. \tag{6.1}$$

Therefore, the source becomes optically thin at around $t \sim 2$ days after the merger (see also Piro & Kollmeier, 2018).

As the shock expands it continues to sweep up surrounding material. At a certain point, the mass of the surrounding material that is swept up becomes larger than the total ejected mass, and the expansion begins to decelerate. This is known as the Sedov-Taylor phase. At that point, the injection of non-thermal particles for acceleration becomes slower and the non-thermal emission should become dimmer. Assuming the remnant expands spherically, therefore occupying a volume $V \sim (\beta ct)^3$, and that the surrounding material has a uniform density of $\bar{n} = 10^4\ \text{cm}^{-3}$, the Sedov-Taylor phase is achieved in a timescale of

$$t = 220\ \frac{0.2}{\beta} \left(\frac{10^4\ \text{cm}^{-3}}{\bar{n}}\right)^{1/3} \left(\frac{M}{10^{-2}\ M_\odot}\right)^{1/3}\ \text{days}. \tag{6.2}$$

In reality, the peak in radio, optical and X-ray fluxes was observed to take place at around 160 days after the merger (van Eerten et al., 2018).

6.2 Magnetic field constraints from radio-to-X-ray observations

We now focus on non-thermal particle acceleration within the source. We analyze the non-thermal spectrum of the remnant as a result of synchrotron emission by a population of accelerated electrons, assuming that this radiation is emitted isotropically by the source. Assuming

[1] The actual value of the outflow speed is not well constrained, so we will assume this value as a benchmark throughout this chapter, and comment on the effect of a different value on the results, whenever relevant.

that electrons in the source are accelerated to a power law through diffusive shock acceleration (DSA, *cf.* Section 2.1.2), we will now show how the observed radio-to-X-ray fluxes can place constraints on the magnetic field strength in the source.

The remnant of GW170817 has been observed to emit a non-thermal spectrum of increasing luminosity both in the radio ($\sim 0.3 - 6 \times 10^{-5}$ eV) and X-ray ($\sim 0.3 - 8 \times 10^{3}$ eV) ranges (Mooley et al., 2018; Ruan et al., 2018). This radiation is consistent with an electron population accelerated to a spectrum $dN/dE_e \propto E_e^{-2}$. We will now show that the lack of a cooling feature in the observed spectrum between radio and X-ray energies can constrain the synchrotron cooling rate taking place in the source.

Dynamical time limited acceleration– A lower limit can be placed on the magnetic field strength in the remnant from a consideration of the acceleration timescale, $t_{\rm acc} = t_{\rm Lar}/\beta^2$, where $t_{\rm Lar} = R_L/c$ is the Larmor period. This relationship corresponds to Eq. (2.1) with an acceleration efficiency given by $\eta = \beta^2$. By requiring that $t_{\rm acc} < t_{\rm dyn}$, this equation converts to $t_{\rm Lar} < 4 \times 10^{-2}\ t_{\rm dyn}$ (considering $\beta = 0.2$, as will be done throughout this study).

Cooling time limited acceleration– The lack of a cooling break in the observed synchrotron spectrum can be used to place a constraint on the maximum magnetic field giving rise to the observed synchrotron emission. The synchrotron cooling time of mono-energetic electrons with Larmor radius giving rise to synchrotron photons with characteristic energy $E_\gamma^{\rm syn}$ is

$$\begin{aligned} t_e^{\rm syn} &= \frac{9}{8\pi\alpha}\left(\frac{m_e}{E_\gamma^{\rm syn}}\right) t_{\rm Lar} \\ &= 2 \times 10^3 \left(\frac{10\ {\rm keV}}{E_\gamma^{\rm syn}}\right) t_{\rm Lar}, \end{aligned} \tag{6.3}$$

which can be derived from Eq. (2.9) by considering the dependence of the Larmor period of the electron on its energy and on the magnetic field strength (see definition of Larmor radius following Eq. (2.1)).

Assuming the electron spectrum is produced via DSA (thus giving rise to a power law spectrum $dN/dE_e \propto E_e^{-2}$, *cf.* Section 2.1.2), the lack of a cooling feature in the observed synchrotron spectrum can then constrain the maximum cooling rate of the electrons at the highest energies, $t_{\rm dyn} < t_e^{\rm syn}(E_\gamma^{\rm syn} = 10\ {\rm keV})$, leading to the constraint $t_{\rm Lar} > 4 \times 10^{-4}\ t_{\rm dyn}$.

We therefore find a dual constraint on the Larmor period of the particles,

$$4 \times 10^{-4} < t_{\rm Lar}/t_{\rm dyn} < 4 \times 10^{-2}, \tag{6.4}$$

where $t_{\rm dyn} \sim 100$ days $\sim 9 \times 10^6$ s.

We know that the energy of the emitted synchrotron photons and the emitting electron energy

are related to the magnetic field via Eq. (2.9). The above constraint on $t_{\rm Lar}/t_{\rm dyn}$ therefore translates into an upper limit on the magnetic field strength in the source.

The dynamical time limit, $t_{\rm Lar} < 4 \times 10^{-2}\ t_{\rm dyn}$, sets the limit on the acceleration energy scale at $\sim 10^{14}$ eV, or a corresponding minimum magnetic field strength of ~ 0.03 mG (by considering the dependence of the Larmor period on the magnetic field and the electron energy as well as Eq. 2.9). Conversely, the cooling time limit $t_{\rm Lar} > 4 \times 10^{-4}\ t_{\rm dyn}$, translated into a minimum electron energy of $\sim 9 \times 10^{12}$ eV and therefore limits the magnetic field strength to a maximum of 2 mG. The overall dual constraint on the magnetic field in the system is therefore given by:

$$0.03 \text{ mG} < B < 2 \text{ mG}. \tag{6.5}$$

In Section 6.4 we will argue that gamma-ray observations can provide a stronger constraint on the minimum magnetic field strength. On the other hand, assuming the maximum value of $B = 2$ mG throughout the age of the remnant, the corresponding maximum energy for protons (not accounting for energy losses) is

$$E_p^{\rm max} \approx 6 \times 10^{15} \left(\frac{\beta}{0.2}\right)^2 \left(\frac{t_{\rm dyn}}{100 \text{ days}}\right) \left(\frac{B}{2 \text{ mG}}\right) \text{ eV}, \tag{6.6}$$

obtained by equating $t_{\rm acc}$ and $t_{\rm dyn}$. While in Section 6.3.2 we will argue for the possibility of a different scenario with strong, gauss-level magnetic fields, it is worth noting that even within the above constraints, this source type is capable of accelerating CRs to energies beyond the knee and below the ankle. This energy region of the CR spectrum is interesting due to the indications for the onset an additional source component (Hillas, 2004), which may help bridge the energy gap between the iron knee (Apel et al., 2011) and the ankle (*cf.* Figure 2.1). Furthermore, observations of a dipole anisotropy in the CR flux at such energies (Giacinti et al., 2012; Abreu et al., 2012) strongly support the idea that this additional component should be extragalactic in origin. This new source class therefore appears as an interesting candidate due to its capability to efficiently accelerate particles.

6.3 Electron synchrotron emission

There are two alternative scenarios that can explain the observed luminosity of the synchrotron emission. With a low magnetic field, like that obtained from the arguments discussed above, an electron energy density is needed that far surpasses the magnetic energy density (see discussion in Section 6.3.1). An alternative scenario with a high magnetic field strength (to boost synchrotron production), would violate the constraints derived previously from the lack of a cooling break

in the spectrum. However, this violation can be negated in two possible ways. Firstly, if the observed synchrotron spectrum is in fact produced by cooled electrons, accelerated in the source to a much harder spectrum than that produced in strong, non-relativistic shocks via DSA. Secondly, if the emitted radiation originates in fresh electrons, picked up at the edge of the outflow and continuously accelerated at the shock front. In the following we discuss possible parameter sets representative of either extreme magnetic field case.

6.3.1 Slow acceleration scenario

In this scenario we consider a weak, uniform magnetic field present in the remnant. As an example, we consider the value $B = 0.2$ mG, which lies within the range given by Eq. (6.5).

As discussed in Section 2.1.2, the peak of the synchrotron power emitted by an electron distribution is dominated by the electrons of highest energies. In this case, we observe X-rays from the source (Troja et al., 2017), from which the total number of highest-energy electrons, $N_e^{\mathrm{X-ray}}$, can be derived. The energy of these electrons is related to the energy of the emitted X-rays through Eq. (2.9), and the total X-ray luminosity is given by Eq. (2.8). This luminosity has been constrained at around 110 days after the merger to a value around $L_{\mathrm{syn}}^{\mathrm{X-ray}} = 4\times10^{39}$ erg s^{-1} at ~ 10 keV (Margutti et al., 2017), yielding $N_e^{\mathrm{X-ray}} = 3 \times 10^{46}$.

Since the energy spectrum of non-thermal electrons goes as $E_e \frac{dN}{dE_e} \propto E_e^{-1}$, the total number of accelerated electrons is dominated by the lowest-energy, radio-emitting electrons, whose number is given by

$$N_e^{\mathrm{radio}} = N_e^{\mathrm{X-ray}} \frac{E_e^{\mathrm{X-ray}}}{E_e^{\mathrm{radio}}}. \tag{6.7}$$

This yields a number of radio-emitting ($E_\gamma = 10^{-6}$ eV) electrons of 3×10^{51}. It is important to note that if the electron population were to extend to energies lower than $E_e^{\mathrm{radio}} \approx 300$ MeV, the total electron population could then reach a number close to the maximum possible number of swept-up and injected electrons at the Sedov-Taylor phase (around $N_e^{\mathrm{total}} = 10^{55}$). This is particularly critical if we realize that electrons are generally thought to carry only 1% of the total energy of non-thermal particles (Burbidge, 1959). It is also informative to note that in this scenario the electron energy density is much higher than the magnetic energy density, $u_e \approx 10^7 u_B$, assuming a volume of 7×10^{50} cm^3 corresponding to a sphere expanding with velocity $\beta = 0.2$ after 110 days.

Finally, note that the maximum energy that may be achieved by electrons in the source may be higher than that which dominates X-ray production. In fact, as shown in the left panel in Figure 6.1, in a 2 mG field electrons can be accelerated in the source up to $E_e^{\mathrm{max}} = 700$ TeV by 110 days. At these energies, acceleration is limited by the age of the remnant and a cutoff is

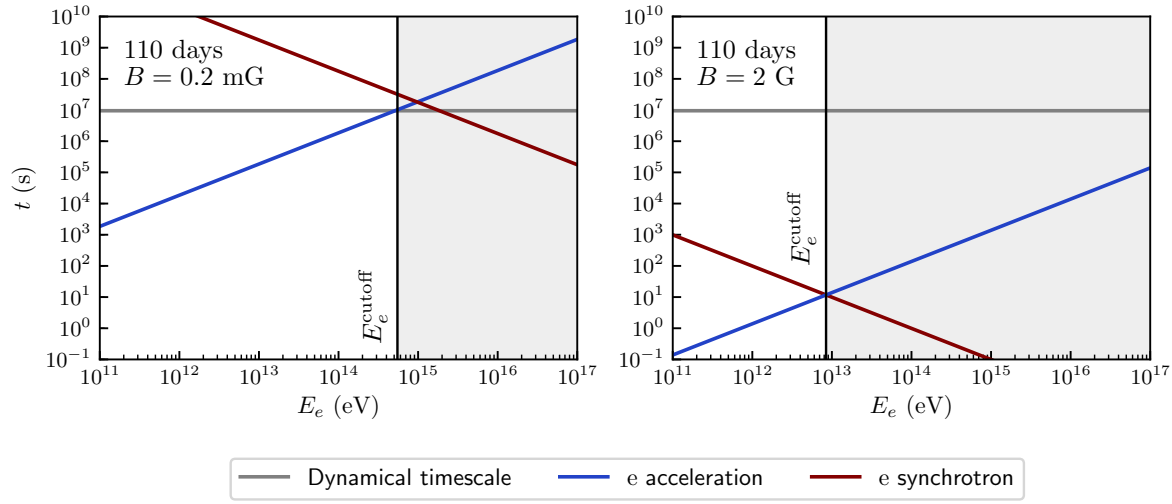

Figure 6.1: Figure taken from Rodrigues et al. (2019), showing the relevant timescales for radiative processes in the remnant of GW170817 for a magnetic field strength of $B = 0.2$ mG *(left)* and $B = 2$ G *(right)*. The dynamical timescale (gray line) corresponds to the expansion time of the source, and determines the efficiency of electron acceleration (blue) and cooling (red). In the case of a weak magnetic field, we expect a cutoff in the electron distribution at around 600 TeV, when acceleration becomes inefficient given the age of the remnant. On the other hand, in the strong magnetic field case, we expect a cooling cutoff at 8 TeV, the energy at which synchrotron cooling dominates over acceleration.

expected. A cooling break is not expected, since synchrotron emission (blue line in Figure 6.1) is not efficient in the relevant energy range.

6.3.2 Fast acceleration scenario

An opposite scenario involving a stronger magnetic field strength is also possible. In fact, a scenario with strong magnetic fields is predicted by CR-driven magnetic field amplification in environments with strong CR fluxes (Bell, 2004). This instability may in principle drive the magnetic field strength up to a saturation value of as much as 2 G $(M/10^{-2}\ M_\odot)^{1/2}(t/100\ \text{days})^{-3/2}$ (assuming $p_{\rm ram} \sim u_{\rm CR}$), an estimate that is independent of the outflow speed.

As mentioned earlier, one possibility to explain observations within a strong magnetic field scenario is that the electron population is accelerated to a harder spectrum by a mechanism other than DSA in strong non-relativistic shocks, such as DSA in mildly relativistic oblique shocks, or even stochastic acceleration (Jones, 1994) (in spite of its low efficiency compared to the fast radiative cooling, which can make it difficult to explain the acceleration of electrons to

high enough energies). The entire non-thermal electron population would then be cooled to a E^{-2} spectrum due to the strong magnetic fields, yielding the observed power law synchrotron emission.

Alternatively, a strong magnetic field may also be explained if the emitted radiation continuously originates from electrons accelerated near the edge of the outflow. In fact, as the remnant expands, the number of electrons picked up grows with the volume, while the magnetic field strength may conceivably decrease linearly with time, as inferred for other outflows like the supernova SN 1993J (Fransson & Bjornsson, 1998; Tatischeff, 2009). This would imply that synchrotron emission is dominated by freshly accelerated electrons, which dominate in number, thereby relaxing the constraint on the synchrotron cooling efficiency of the source.

Regardless of the origin of the synchrotron emission, in a strong magnetic field scenario the observation of radio emission can be used to constrain the efficiency of synchrotron self-absorption in the remnant at those energies (*cf.* Section 6.4), which in turn can be used to constrain the magnetic field to a maximum of 10 G.

For such strong magnetic fields, the observed X-ray luminosity is produced by lower energy electrons than in the weak magnetic field scenario. Adopting the value $B = 2$ G, we have $E_e^{\mathrm{X-ray}} \approx 300$ GeV (*cf.* Eq. 2.9). To account for the observed X-ray luminosity, the necessary number of X-ray-emitting electrons is $N_e^{\mathrm{X-ray}} = 3 \times 10^{42}$, which corresponds to a number of radio-emitting electrons of 3×10^{47} (Eq. 6.7), a value much lower than that derived for the weak magnetic field scenario. In this case, the energy density of the non-thermal electrons and magnetic field are related by $u_e \approx 10^{-7} u_B$, which means the outflow is strongly magnetically dominated.

Finally, as shown in the right panel of Figure 6.1, the maximum electron energy allowed in this case is only 8 TeV. At that energy synchrotron cooling dominates over acceleration at higher energies, and a cooling cutoff is expected (Aharonian, 2000), a characteristic that is not expected in the low magnetic field scenario.

6.4 Leptonic modeling of the multi-wavelength emission

A population of non-thermal particles embedded in ambient radiation fields invariably give rise to subsequent high-energy emission from inverse Compton scattering. Assuming that synchrotron emission dominates the radiation field seen by the non-thermal electrons, then the inverse Compton emission produced will be dominated by synchrotron self-Compton scattering (SSC). As discussed in Section 2.4.1, considerable SSC emission can be expected for cases in which $u_e \gg u_B$ (Fransson & Björnsson, 2005), which means that a weak magnetic field scenario may be possible to constrain by observational gamma-ray flux limits.

In order to determine the inverse Compton emission at different times, a prescription for the radiation field evolutions must be adopted. The non-thermal radiation field is normalized by fixing the X-ray luminosity, obtained from observations at 9, 15 and 110 days after the merger (Margutti et al., 2017; Troja et al., 2017; Ruan et al., 2018), and subsequently extrapolating it back in time assuming a continuation of the inferred power-law evolution of the form,

$$L_{\mathrm{X-ray}} = 2 \times 10^{39}\ (t/110\ \mathrm{days})^{0.6}\ \mathrm{erg\ s^{-1}}. \tag{6.8}$$

At early times, $t < 15$ d, a bump is observed in the optical-to-infrared range of the SED, with a spectral shape characteristic of thermal emission (Villar et al., 2018). This thermal bump is seen to become dimmer between 10 and 74 days after the event (Villar et al., 2018), at which time it becomes out-shined by the non-thermal emission. Based on these observations, we model the evolution of this thermal luminosity as

$$\begin{aligned} L_{\mathrm{th}} &= 4 \times 10^{40}\ \mathrm{erg\ s^{-1}} && (t <\ 7\ \mathrm{days}) \\ &= 4 \times 10^{40}\ (t/7\ \mathrm{days})^{-2.3}\ \mathrm{erg\ s^{-1}} && (t > 7\ \mathrm{days}). \end{aligned} \tag{6.9}$$

In order to estimate the emission spectrum, I developed a semi-analytical model for leptonic interactions in the merger remnant. The processes included are synchrotron emission and inverse Compton scattering (Blumenthal & Gould, 1970), synchrotron self-absorption (Longair, 1994) and photon-photon annihilation (Gould & Schreder, 1967).

In Figure 6.2 we show the resultant broadband SED produced by the remnant at 5 and 110 days after the merger. In the left panel we show the case where a 2 mG magnetic field is present in the outflow at 5 days after the merger event. At these early times in the remnant, the thermal radiation field (red curve) provides the dominant target for inverse Compton emission (yellow curve). We also show the H.E.S.S. upper limit, in the range $0.5 - 6$ TeV, at 5.2 days (Abdalla et al., 2017). The magnetic field at these early timescales has been adopted sufficiently high so as to ensure that the inverse Compton emission does not overshoot the H.E.S.S. upper limit. For this case, a sharp cutoff is introduced into the inverse Compton spectrum, due to pair production on the thermal radiation field. Note also the low-energy cut-on of the synchrotron spectra (blue curves) is given by synchrotron self-absorption, as discussed in Section 2.4.1. The considerable level of inverse Compton emission found for the case of low magnetic field demonstrates that at later times, once the thermal bump has reduced sufficiently, a lower limit on the magnetic field strength may be placed at late timescales, by follow-up TeV observations of the remnant. We do not account for the cascades that would result from pair production, since the effect on the spectrum would impact low energies that may be neglected for our purpose of comparing the

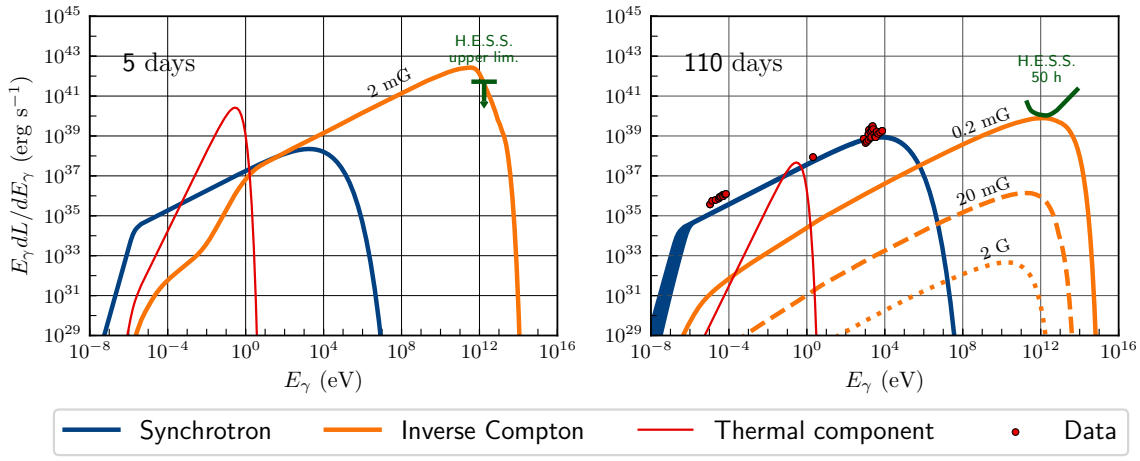

Figure 6.2: Luminosity spectrum produced by electron synchrotron emission (blue, including synchrotron self-absorption), and inverse-Compton scattering (yellow), at five days (left) and 110 days (right) after the merger event. For the 110 day case we show the results for three values of B-field strength. In red we show optical (Lyman et al., 2018), radio and X-ray data (Margutti et al., 2018) at 110 days. In the left panel we show in green the limit on the TeV luminosity at 5 days (Abdalla et al., 2017) and on the right panel we show the H.E.S.S. sensitivity (Holler et al., 2016), which translates into an upper limit on the source gamma-ray luminosity. Figure taken from Rodrigues et al. (2019).

inverse Compton flux to the H.E.S.S. observation limit.

In the right panel of Figure 6.2 we show the resultant broadband SED produced by the remnant at 110 days, normalized to the observed X-ray luminosity, and we show the resultant inverse Compton emission for three values of magnetic field strength. At this later time in the remnant, the synchrotron radiation field provides the dominant target for inverse Compton emission. The red points in this plot show radio, optical, and X-ray data points taken at this time (Margutti et al., 2018; Lyman et al., 2018).

In the weak magnetic field scenario ($B = 0.2$ mG, right panel of Figure 6.2), the inverse Compton luminosity dominance is highest and the predicted gamma-ray luminosity is 10^{40} erg s^{-1}, peaking at $E_\gamma = 1$ TeV. We can see that this emission is at the level of the H.E.S.S. 50 h sensitivity (green curve, Holler et al., 2016)), which indicates that this instrument has the required sensitivity to set a lower limit on the magnetic field strength in the source. In contrast, for higher magnetic field strengths, we see that the low electron density yields a gamma-ray power of only 3×10^{32} erg s^{-1} with a 10 GeV peak. This result demonstrates that the potential probing of the Compton peak by TeV instruments can provide fresh insights to discriminate between the

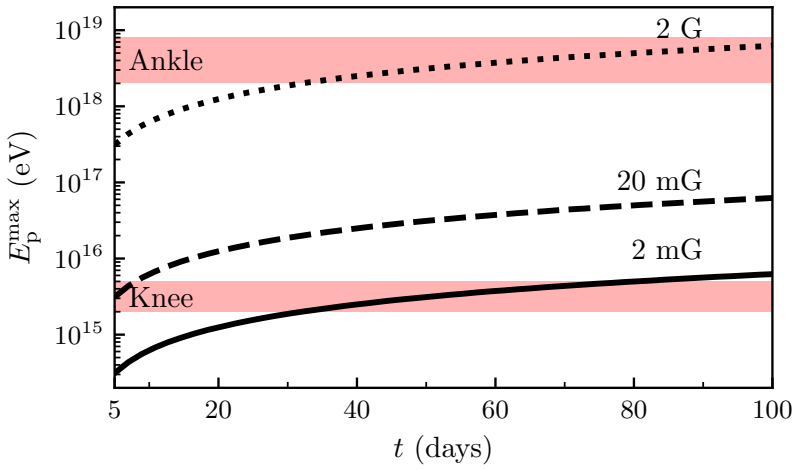

Figure 6.3: Maximum energy achieved by protons accelerated in the remnant as a function of time since the merger, for three different magnetic field strengths. The age of the remnant is always the limiting factor to the maximum proton energy, as photo-meson production is never efficient. As throughout this work, an expansion velocity of $\beta = 0.2$ has been assumed. Figure taken from Rodrigues et al. (2019).

different magnetic field strengths considered.

6.5 Acceleration of cosmic-ray nuclei

Besides the observational evidence for a non-thermal population of electrons in the remnant, the mildly relativistic speeds of the outflow also indicates it to be a potential CR acceleration site, based on purely theoretical arguments (Bell et al., 2018). These facts motivate the study of the acceleration and radiative interactions of CR nuclei in the source, which will be addressed in this section.

6.5.1 Acceleration to energies between the knee and the ankle

The determination of the interaction processes which dictate the maximum energy of the accelerated nuclei depends on the magnetic field strength adopted. Since the low magnetic field value discussed in Section 6.3.1 would increase the acceleration time, acceleration up to high energies and the onset of nuclear photo-disintegration would not be possible. For the purpose of investigating these interactions, we adopt in the following the extreme magnetic field scenario discussed in Section 6.3.2.

As shown in the following section, photo-hadronic cooling of protons in the source is never efficient, even with the bright optical emission at early days. Instead, the maximum energy achieved by protons is always limited by the age of the remnant. In Figure 6.3 we show the evolution of the maximum proton energy with time, for three constant values of magnetic field strength. We can see that in a weak magnetic field scenario, the source can only accelerate CRs above the knee from around 80 days after the merger onwards. On the other hand, in the case of a strong magnetic field scenario, late-time acceleration of protons is possible up to the ankle, as shown in Figure 6.3 for the maximum magnetic field strength (2 G) we consider.

6.5.2 Energy losses at the ankle

Like electrons, nuclei accelerated by the source will also interact with the target photons present within it. In fact, the appreciable attenuation at early times (< 10 days) of TeV photons in the source found in Section 6.4 gives reason to expect also considerable photo-disintegration in the source on these timescales (Neronov et al., 2009; Murase & Beacom, 2010; Aharonian & Taylor, 2010).

In Figure 6.4 we show the interaction timescales of the different processes at work for both protons (left panel) and iron-56 nuclei (right panel), obtained using the NeuCosmA code, also applied in Chapters 3 and 4. We take as reference the time interval of 9 days after the merger, when photo-hadronic interactions are most efficient, as discussed later in this section. In the case of protons, we see that acceleration (red line) is always limited by the age of the remnant, which corresponds to the dynamical timescale of the system, shown in the gray line. On the other hand, the source is seen to be optically thin to photo-meson production (blue curve). We also plot the pair production loss time (magenta curve), showing that this process is always sub-dominant as an energy loss mechanism for hadrons.

On the right panel of Figure 6.4, we can see that at the 9 day timescale the bright optical radiation can efficiently photo-disintegrate iron-56 nuclei, as shown by the yellow curve, achieving an optical thickness (i.e. the ratio between the dynamical and photo-disintegration timescales) of about 10 at 2 EeV. At this energy, iron nuclei cannot be efficiently accelerated, as photo-disintegration becomes the dominant process and the nuclei instead disintegrate into lighter isotopes. At later times, however, the thermal photon luminosity decreases (Eq. 6.9), making photo-disintegration less efficient, and the maximum energy becomes limited only by the age of the remnant, as in the case of protons.

In Figure 6.5 we show the temporal evolution of the optical depth of the source to photo-disintegration of different isotopes, as well as photo-meson production by protons. Here, the optical depth is defined as $\tau_{\rm int} = t_{\rm dyn}/t_{\rm int}|_{\rm E_{max}}$, where $t_{\rm dyn}$ is the dynamical timescale and $t_{\rm int}$

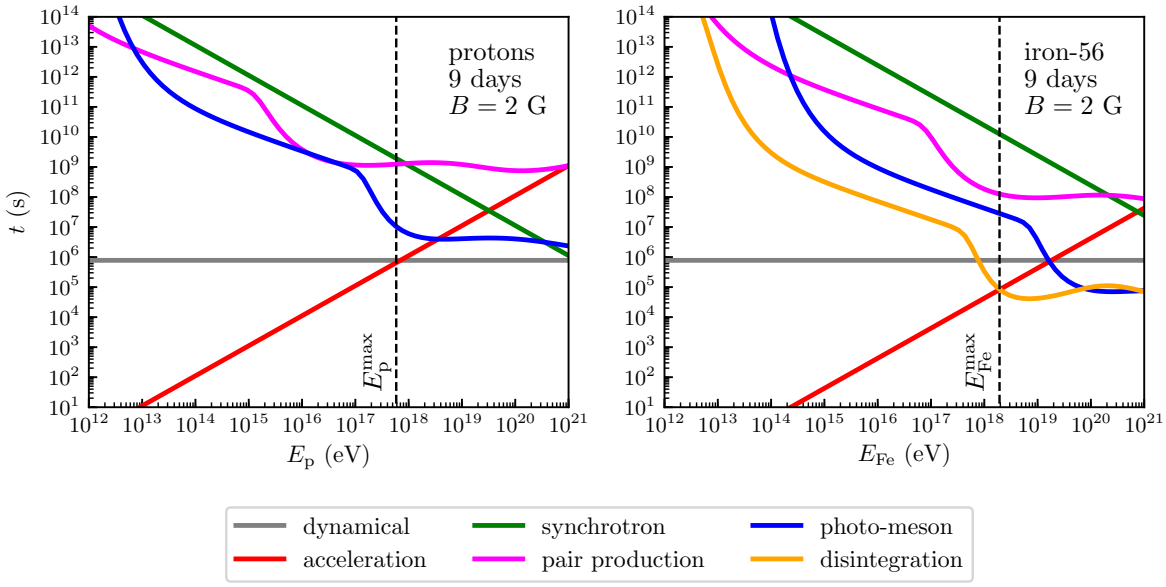

Figure 6.4: Interaction times for protons (left) and iron-56 nuclei (right) as a function of the particle energy, 9 days after the merger, in the high magnetic field scenario. Note that for the synchrotron and pair production processes we show the energy loss timescale. Figure taken from Rodrigues et al. (2019).

the interaction timescale, evaluated at the maximum energy of the CR (see vertical dashed line in Figure 6.4). This provides a measure for the fraction of accelerated CRs that escape the source. The optimal time for nuclear photo-disintegration is found to be reached within the first weeks after the merger. This happens because of the competing processes within the expanding object, namely the rise of the maximum particle energies achievable and the decrease of the thermal photon field density.

Assuming continuous injection, the late time acceleration of protons can bring them to energies up to the ankle. On the other hand, the composition of CRs eventually escaping the source can be quite different depending on the period of efficient photo-disintegration, on the primary isotope, and on the escape mechanism assumed. As the present study is held very general, we do not make an effort to calculate the resultant late-time accelerated nuclei spectra.

6.5.3 Cosmic rays from a merger population

With some level of photo-disintegration expected within the source during early times, and very high energies becoming within reach at later times, we finally turn our attention to the ensemble population of such sources. We implicitly assume here that the event we discuss throughout

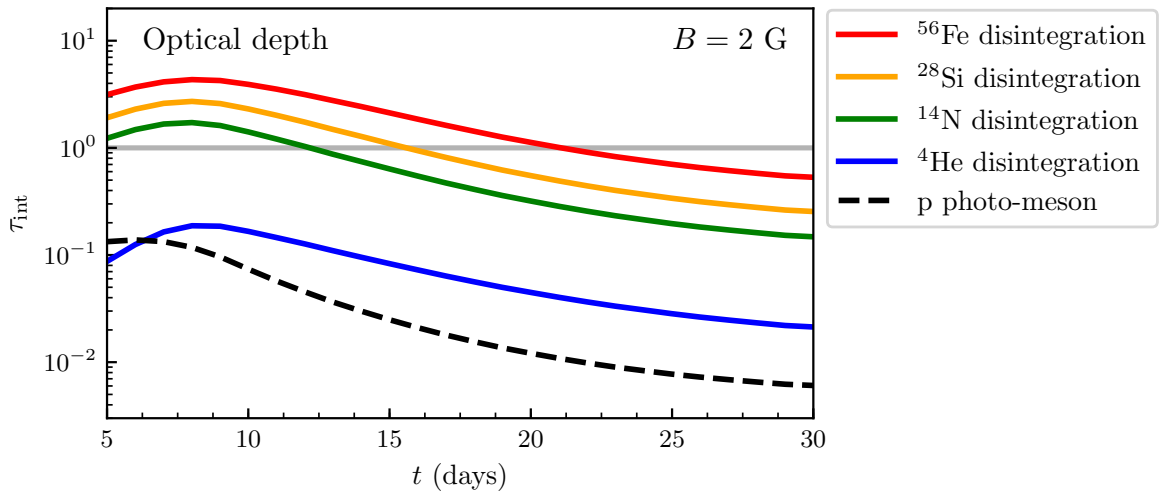

Figure 6.5: Optical depth of the remnant to photo-disintegration of different isotopes (solid curves) and photo-meson production by protons (dashed curve) as a function of the time since the merger, in the high magnetic field scenario. The horizontal gray line represents the transition from optically thin to thick (see main text). Figure taken from Rodrigues et al. (2019).

this work is representative of a population of identical sources which could accelerate CR nuclei. In this scenario, there should be a number of electromagnetic counterparts to this population, with the event EM170817 being the only one detected so far.

To support a CR spectrum at Earth with the abundance observed (Gaisser et al., 2013), a local emissivity $\mathcal{L}_0$ of their sources is required. Assuming a CR luminosity density of $\mathcal{L}_0 \sim 4 \times 10^{44}$ erg yr^{-1} Mpc^{-3} (Waxman, 1995), and adopting a specific local rate for these sources, one can infer the required energy released per event. Considering a local rate of neutron star mergers of 1540^{+3200}_{-1220} Gpc^{-3} yr^{-1} (Abbott et al., 2017a), we estimate the required energy input into CRs in each merger event of $E_{\rm CR} \approx 2 \times 10^{50}$ erg. This estimate is roughly consistent with a fraction of $\sim$10% of the released outflowing kinetic energy given in Section 6.1.

Using the previous value for the total energy output in CRs, we can also estimate the total neutrino fluence produced by the source. From Figs. 6.4 and 6.5 we see that the optical thickness of the source at the 9 day timescale is about $\tau_{\rm p\gamma} = 0.1$, while for later times it diminishes due to the dimming of the thermal luminosity (Eq. 6.9) and the overall reduction in photon density due to expansion. The average optical thickness from 5 to 160 days is then approximately $\langle \tau_{\rm p\gamma} \rangle = 5 \times 10^{-3}$. The resulting estimate of the neutrino fluence from the source is $E_{\rm CR} K_{p\gamma} \langle \tau_{\rm p\gamma} \rangle / (4\pi d^2) \approx 5 \times 10^{-4}$ GeV cm^{-2}, where $K_{p\gamma} \approx 0.15$ describes the fraction of energy taken by a pion in each photo-meson production interaction and $d = 40$ Mpc is the distance to the source. Note that the normalization factor for this estimate is provided by the energy input in CRs deduced above,

which carries considerable uncertainty, mainly due to the present large uncertainty on the source rate. However, a simple comparison with the present upper limits for the neutrino flux from EM170817 (Albert et al., 2017) indicates that the level of this flux would be challenging to reach by present instruments like IceCube. For comparison, Fang & Metzger (2017); Kimura et al. (2017) considered the case of neutrino emission from pulsars formed after neutron star merger events. In both cases, rather increased neutrino fluxes were found to be expected compared to those determined here, which is due to the faster acceleration process at play in the source and the differing source environment.

6.6 Summary

In this chapter, we have considered non-thermal particle energy losses within the fast-moving remnant outflow associated with EM170817. Assuming that this emission originates from electrons, the consideration of the lack of a cooling break in the synchrotron emission from this remnant, whose age is accurately known, allowed a constraint to be made on the strength of the source magnetic field at the mG level. The subsequent synchrotron self-Compton emission expected demonstrated that for such a magnetic field strength level, large inverse Compton emission is expected. Current gamma-ray sensitivities are shown to be able to constrain source properties like the magnetic field strength.

An alternative strong, gauss-level magnetic field scenario was also proposed, in which two potential origins of the synchrotron emission are discussed. The first is that it is produced by electrons with a hard injection spectrum, which subsequently cool to an E^{-2} distribution. The second possibility is that this synchrotron emission is constantly produced by fresh electrons, picked up and accelerated near the edge of the remnant, as inferred for other similar astrophysical outflows.

For the strong magnetic field scenario considered, we also tested the efficiency of photo-meson production and photo-disintegration of CRs. The thermal component was demonstrated to be crucial for the interactions of the CRs in the outflow of the binary neutron star merger remnant, especially at early times, when it provides the dominant target radiation field for these processes. While for the assumed outflow speed the source is always optically thin to photo-meson production, photo-disintegration of some nuclear isotopes is efficient in the first ~ 10 days after the merger event. Later, when the volume of the object has expanded, the thermal component becomes dimmer and the system is optically thin to all hadronic interactions. At this point, the maximum energy of nuclei is no longer dictated by losses on this thermal radiation, allowing the acceleration up to energies beyond the ankle, while protons can be accelerated up to the ankle. In order for such a population to power the observed CR flux, a total energy

output in CRs by the source of $\sim 10^{50}$ erg is required at late timescales.

An origin of CRs below the ankle related to neutron star merger remnants carries similarities with ideas put forward in other works (Anchordoqui et al., 2008; Aloisio et al., 2014; Globus et al., 2015; Unger et al., 2015; Aab et al., 2017b; Biehl et al., 2017). Indeed, such a scenario can naturally explain the apparent lightness of the CR composition inferred at these energies from elongation rate measurements by Auger (Aab et al., 2017c). The recent in-depth observations of the remnant of GW170817, over a wide energy range, have shed light on the non-thermal aspects of this type of phenomena, and motivates them as promising CR sources.

Chapter 7

Conclusion

The sites of acceleration of the UHECRs are not yet known. The observation by IceCube of neutrinos with energies up to several PeV seem to imply that some of the sources where UHECRs are accelerated, are sited of hadronic interactions leading to high-energy neutrino emission from the decay of secondaries. However, the neutrino data collected thus far does not seem to exhibit any statistically relevant overlap with the high-energy sources observed in gamma rays. On the one hand this fact is surprising, given that high-energy gamma rays are co-produced with neutrinos in photo-meson interactions of UHECRs. On the other hand, we know that neutrinos can travel cosmological distances unscathed, while high-energy photons cascade to lower energies during propagation due to interactions with the CMB, and are therefore limited to the local Universe. It may therefore be possible that the sources of neutrinos are simply too dim or distant to be observed by gamma-ray telescopes. Nonetheless, the study of neutrino production in CR accelerators is a valuable tool for the understanding of the high-energy Universe. The majority of this work was dedicated to the study of blazars as sources of UHECRs and neutrinos. Blazars are characterized by the emission of a multi-wavelength spectrum ranging from radio to gamma rays; this emission is strongly beamed due to the relativistic speed of the jet, which is pointing directly at the observer. Moreover, blazars are observed to undergo flares, which are periods of more intense emission, often accompanied by fast variability of the emitted flux, especially in gamma rays. This behavior can in general be explained through a variety of models that include the acceleration of electrons and/or nuclei to a non-thermal spectrum.

In this work we explored the paradigm of lepto-hadronic models, where both electrons and nuclei are accelerated in plasma shocks in the jet. This type of model has the advantage that it can be used to explain simultaneously the emission of photons, CRs and neutrinos. Although there is now evidence for the presence of nuclei heavier than protons in the observed UHECR spectrum, most current lepto-hadronic blazar models focus mainly on the case where only protons and electrons are accelerated in the jet. For the first time, we have computed numerically the effect of the acceleration of different isotopes on the emitted neutrino spectrum. For that, we used a one-zone blazar model (Rodrigues et al., 2018a) based on a numerical solver (the

NeuCosmA code, Baerwald et al., 2013) to simulate the evolution of the particle densities in the jet. The simulation accounts for all relevant photo-hadronic and electromagnetic processes, including all channels of production of secondary nuclear species. We then applied this model to blazars of different luminosities. The photon spectrum in the jet, considered static, was obtained from observations by following the blazar sequence (Fossati et al., 1998; Ghisellini et al., 2017), which provides an average spectrum for each value of gamma-ray luminosity, based on a large sample of observed blazars.

Regarding nuclear interactions in blazar jets, we find two distinct regimes, depending on the blazar luminosity. In low-luminosity sources, the injected isotope does not interact efficiently in the source, due to the low density of photons in the jet, and therefore the CR spectrum emitted by the blazar has the same chemical composition as the accelerated CRs. Assuming that the strength of the magnetic field in blazar jets scales with the emitted luminosity, then low-luminosity blazars have less capability of accelerating CRs to ultra-high energies; however, because of the low density of the photon fields in the jet, the accelerated CRs can efficiently escape before cooling through interactions with this photon field. In other words, low-luminosity blazars are shown to be optically thin to photo-hadronic and nuclear interactions. Therefore, the maximum energy of the ejected CR spectrum is the same as that of the accelerated spectrum. This energy scales linearly with the charge of the accelerated CRs (and therefore with the number of protons of the accelerated isotope).

High-luminosity blazars, on the other hand, can accelerate CRs to higher energies due to the stronger magnetic field strength in the jet. However, a photon field density high enough will lead to efficient interactions of the CRs at the highest energies, thereby decreasing the maximum energy of the CR spectrum that is effectively emitted by the source. This means that very bright blazars, like FSRQs, are generally inefficient emitters of UHECRs. The most efficient UHECR emitters in the blazar family are therefore sources of intermediate luminosity, such as ISP BL Lacs (Figure 3.11). The maximum energy of CRs emitted by these intermediate-luminosity blazars was shown to have an order of magnitude of ~ 10 EeV, both for iron-56 and for protons. Note that this result was obtained for the ideal scenario of ultra-efficient acceleration (*i.e.* an acceleration efficiency of $\eta = 1$); in reality, particle acceleration in relativistic shocks is though to be inefficient (*e.g.* Araudo et al., 2015, 2016), which shows clearly that the emission of UHECRs is challenging from a theoretical standpoint.

While bright blazars are inefficient CR emitters, they can efficiently produce neutrinos through interactions of the accelerated nuclei with the dense photon fields in the source, while low-luminosity blazars have low neutrino production efficiency (Figure 3.12). Neutrinos from high-luminosity blazars originate mainly in the decay of pions from photo-meson interactions. Beta decay of neutrons and unstable secondary nuclei also leads to the emission of neutrinos of lower

energy, but this contribution is sub-dominant in all the cases studied.

In the case of FSRQs, a type of high-luminosity blazars, this high neutrino production efficiency is made even higher by the external photon fields produced in the BLR, consisting of thermal emission and atomic broad lines. The contribution of the external fields to neutrino production is two-fold: firstly, a part of these photon fields will appear relativistically boosted in the rest frame of the jet, and is therefore seen by the accelerated CRs as an additional target field for photo-hadronic interactions. Secondly, the CRs that escape the jet have to transverse the BLR itself, as well as the larger radiation field of the dusty torus, before escaping into the intergalactic medium. To account for this, we considered a three-zone model for the modeling of high-luminosity FSRQs. We showed that the contribution of the external radiation zones to the overall neutrino flux can contribute at an equal level as that from the jet, or even dominate at the peak of the spectrum. This is especially so in the case of fast CR escape from the jet, such as through advection by plasma winds.

The fact that high-energy blazars are more efficient neutrino emitters seems to contradict the lack of coincidences between the observed neutrino arrival directions and known gamma-ray sources. In order to asses this more general problem, we applied the model developed previously to the cosmological distribution of blazars, as parametrized by Ajello et al. (2012, 2014) based on the observed sample of gamma-ray sources. We then tested the possibility that the diffuse flux observed by IceCube is emitted by this blazar distribution.

As mentioned previously, the non-observation of coincidences between neutrinos and gamma-ray blazars is a major challenge to this hypothesis, since it places an upper limit on the neutrino fluxes emitted by observed sources (so-called *stacking limit*). However, cosmological evolution models predict that there is a large number of dim blazars, concentrated mostly at redshift $z < 1$, most of which are of the BL Lac type (*i.e.*they not exhibit the broad line emission that is typical of FSRQs). These BL Lacs emit an individual flux in gamma rays that is below the detection threshold of the Fermi LAT, and are therefore not observable in this wavelength. We have shown that if this vast population of dim blazars has a baryonic loading (roughly the ratio of proton-to-electron energy density) that is much larger than that of high-luminosity sources, the entire flux of sub-PeV to PeV neutrinos observed by IceCube can be well explained by blazar emission (Figure 4.7). In detail, we have concluded that the baryonic loading of blazars should scale as a power law of their gamma-ray luminosity (Palladino et al., 2019). The important fact is that in order to explain the totality of the diffuse neutrino spectrum with blazars, the contribution of bright sources must be suppressed in order to circumvent the stacking limit. On the contrary, other assumptions that are more commonly adopted and possibly more natural, such as a linear scaling between the neutrino and gamma-ray luminosity of blazars (*e.g.* Kadler et al., 2016; Halzen & Kheirandish, 2016; Wang & Li, 2016; Righi et al., 2017), do not allow for

the blazar origin of the IceCube neutrinos, since the stacking limit would be violated.

This work also addressed what are possibly the two most celebrated discoveries in the last two years in the field of multi-messenger astrophysics: the observation in 2017 of a neutrino in directional and temporal coincidence with a flare of blazar TXS 0506+056 (Aartsen et al., 2018b), and the first coincident detection, earlier the same year, of a gravitational wave signal (GW170817) and electromagnetic emission originating in the same phenomenon (Abbott et al., 2017a). Besides being the first multi-messenger observation including gravitational waves, GW170817 also represents the first ever observation of the merger of a binary neutron star system.

The discovery of the first neutrino blazar is particularly relevant in the context of the main discussion of this work; firstly, because the multi-wavelength emission of the source during the flare provides strong constraints that allow to directly test neutrino production models. On the other hand, it presents itself as the first exception to the lack of correlations between neutrino directions and source positions, therefore raising a number of questions regarding neutrino emission from the entire blazar population. Indeed, based on our blazar model, we demonstrated in Chapter 4 that there are hundreds of blazars currently resolved in gamma rays capable of emitting at least as many neutrinos as TXS (Palladino et al., 2019, see Figure 4.8), which leads to the question of why only one coincident observation has been made so far. A natural assumption would be that TXS might be an outlier and have a higher neutrino production efficiency compared to other blazars with similar luminosities and redshift. In fact, this might be supported by recent findings that TXS may be a masquerading BL Lac (Padovani et al., 2019), which would mean it possesses external fields that boost neutrino production. On the other hand, as discussed by Strotjohann et al. (2018), it may instead be the case that TXS belongs to a large population of similar sources, each with a lower individual neutrino output, and it happened to be the first of those sources to be detected.

In Chapter 5, we based ourselves in some of the ideas put forth in recent works about the 2017 TXS flare (see references in Section 2.5), in order to test whether the neutrino signal observed in 2014-15 from the direction of TXS could have originated in photo-hadronic interactions in the same source (Rodrigues et al., 2018b). Aartsen et al. (2018c) have claimed that the signal consists of 13 ± 5 events above background, with a confidence level of $3.5\,\sigma$. The main difference between this observation and the 2017 flare is that the blazar did not display any enhanced electromagnetic activity, which poses a challenge to photo-hadronic models and seems to potentially indicate a different nature of the two phenomena.

An extensive scan of the physical parameters was performed, using optimization methods such as a genetic algorithm, in order to search for the conditions that would explain the highest number of neutrinos while obeying the multi-wavelength constraints. We also tested different photo-hadronic models with varying degrees of complexity regarding the geometry of the source.

We concluded that at most 5 events in IceCube can be explained, by considering an external field model similar to the high-luminosity FSRQ model discussed previously. This model is particularly well motivated by the recent indications that TXS may in fact possess external fields characteristic of FSRQs (Padovani et al., 2019). However, the low number of predicted IceCube events is incompatible with the observed signal at a 1.6 σ level. Other simpler scenarios, like a one-zone model and a two-zone model involving a more dense radiation zone within a larger blob, predict an even smaller number of events. This study makes clear the challenges involved in explaining efficient neutrino production in TXS in 2014-15, during a quiescent state. Moreover, it shows the importance of multi-wavelength data, especially X-rays, for constraining models of neutrino emission, due to the extensive cascades expected in the source, a conclusion that is supported by the recent results of Reimer et al. (2018).

It is worth noting that if the 2014-15 neutrino signal did indeed originate in TXS, the processes involved must have been substantially different from those behind the 2017 flare, given the different characteristics of the electromagnetic emission during both episodes. From a modeling perspective, this raises the additional challenge of explaining a natural transition between the two states. For instance, the presence of a radiatively efficient accretion disk would be expected to contribute continuously to the attenuation of multi-TeV gamma rays, while in the 2017 flare no such attenuation feature is observed in the source spectrum (*cf.* Figure 2.11).

Finally, the study of the remnant of the merger event GW170817 has shown that this newly-discovered class of extragalactic objects is a promising accelerator of CRs between the knee and the ankle (Rodrigues et al., 2019). This study was possible due to the multi-wavelength follow-up campaign of the remnant, whose data was used as input for a particle interaction model. The fluxes detected in the range from radio to X-ray can be well explained by synchrotron emission from an electron population accelerated in the source through DSA. If that is the case, then the same data can constrain the magnetic field strength in the source from simple considerations regarding the energetics of synchrotron emission. However, we showed that the non-observation of the source in gamma rays by H.E.S.S can provide an even stronger constraint on the minimum magnetic field strength. This result was obtained by means of a semi-analytic model of the electron interactions in the source (namely a synchrotron self-Compton model).

During the first $\sim$ 10 days upon the merger event, the remnant was shown to be optically thick to photo-disintegration of any elements heavier than nitrogen. This is due to the bright kilonova that took place in this timescale, which then quickly subsided. On the contrary, at late timescales ($\gtrsim$ 100 days after the merger), the remnant was shown to be capable of accelerating CRs up to a maximum energy lying anywhere between the knee of the CR spectrum (assuming a conservative value of $B = 2$ mG) and the ankle (with the more aggressive assumption of gauss-level magnetic fields). We have also shown that neutrinos are not expected to be emitted

efficiently, since the source is optically thin to photo-meson production, even at early timescales when the radiation field density is highest.

In general, this thesis addressed the problem of UHECR interactions, and our results contribute to a better understanding of the spectrum and composition of CRs emitted by high-energy sources. The future research can exploit the findings presented in this work in several possible directions, a few of which I briefly outline here. Firstly, our model of nuclear interactions in blazars utilizes the observed multi-wavelength data as an input for the numerical calculation of the interactions. However, in order to apply such a model to particular sources, it is necessary to explicitly calculate the photons emitted from these interactions, so that the blazar SED can effectively constrain the model. Regarding CR emission from the entire blazar population, the next necessary step is to combine the photo-hadronic blazar model with a numerical calculation of CR propagation in the intergalactic medium (similarly to the model applied by Biehl et al. (2017) to the study of GRB emission). Such a paradigm should make it possible to constrain the properties of hadronic blazars based on the observed spectrum and composition of UHECRs. Both these directions are currently being pursued in new projects at DESY Zeuthen, and will hopefully further our understanding of this rapidly developing field. Importantly, our results on both blazars and neutron star merger remnants highlight the fact that continued progress in this field will rely heavily on the synergy between the theory sector and present and future multi-messenger experiments.

Bibliography

Aab, A., et al. 2015a, The Pierre Auger Cosmic Ray Observatory, Nucl. Instrum. Meth., A798, 172. arXiv:1502.01323

—. 2015b, Improved limit to the diffuse flux of ultrahigh energy neutrinos from the Pierre Auger Observatory, Phys. Rev., D91, 092008. arXiv:1504.05397

—. 2016, Testing Hadronic Interactions at Ultrahigh Energies with Air Showers Measured by the Pierre Auger Observatory, Phys. Rev. Lett., 117, 192001. arXiv:1610.08509

—. 2017a, Observation of a Large-scale Anisotropy in the Arrival Directions of Cosmic Rays above 8×10^{18} eV, Science, 357, 1266. arXiv:1709.07321

—. 2017b, Combined fit of spectrum and composition data as measured by the Pierre Auger Observatory, JCAP, 1704, 038. arXiv:1612.07155

—. 2017c, Inferences on mass composition and tests of hadronic interactions from 0.3 to 100 EeV using the water-Cherenkov detectors of the Pierre Auger Observatory, Phys. Rev., D96, 122003. arXiv:1710.07249

Aartsen, M. G., et al. 2013, Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector, Science, 342, 1242856. arXiv:1311.5238

—. 2014, Energy Reconstruction Methods in the IceCube Neutrino Telescope, JINST, 9, P03009. arXiv:1311.4767

—. 2015, A combined maximum-likelihood analysis of the high-energy astrophysical neutrino flux measured with IceCube, Astrophys. J., 809, 98. arXiv:1507.03991

—. 2016, Observation and Characterization of a Cosmic Muon Neutrino Flux from the Northern Hemisphere using six years of IceCube data, Astrophys. J., 833, 3. arXiv:1607.08006

—. 2017a, The contribution of Fermi-2LAC blazars to the diffuse TeV-PeV neutrino flux, Astrophys. J., 835, 45. arXiv:1611.03874

—. 2017b, Extending the search for muon neutrinos coincident with gamma-ray bursts in IceCube data, Astrophys. J., 843, 112. arXiv:1702.06868

—. 2017c, The IceCube Realtime Alert System, Astropart. Phys., 92, 30. arXiv:1612.06028

—. 2017d, The IceCube Neutrino Observatory - Contributions to ICRC 2017 Part II: Properties of the Atmospheric and Astrophysical Neutrino Flux. arXiv:1710.01191

—. 2018a, Search for steady point-like sources in the astrophysical muon neutrino flux with 8 years of IceCube data. arXiv:1811.07979

—. 2018b, Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A, Science, 361, eaat1378. arXiv:1807.08816

—. 2018c, Neutrino emission from the direction of the blazar TXS 0506+056 prior to the IceCube-170922A alert, Science, 361, 147. arXiv:1807.08794

—. 2018d, 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, 062003. arXiv:1807.01820

—. 2019, Investigation of two Fermi-LAT gamma-ray blazars coincident with high-energy neutrinos detected by IceCube. arXiv:1901.10806

Abbasi, R. U., et al. 2017, Search for EeV Protons of Galactic Origin, Astropart. Phys., 86, 21. arXiv:1608.06306

Abbott, B., et al. 2017a, GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral, Phys. Rev. Lett., 119, 161101. arXiv:1710.05832

Abbott, B. P., et al. 2016, Observation of Gravitational Waves from a Binary Black Hole Merger, Phys. Rev. Lett., 116, 061102. arXiv:1602.03837

—. 2017b, Gravitational Waves and Gamma-rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A, Astrophys. J., 848, L13. arXiv:1710.05834

Abdalla, H., et al. 2017, TeV gamma-ray observations of the binary neutron star merger GW170817 with H.E.S.S, Astrophys. J., 850, L22. arXiv:1710.05862

Abdo, A. A. 2010, The First Catalog of Active Galactic Nuclei Detected by the Fermi Large Area Telescope, Astrophys. J., 715, 429. arXiv:1002.0150

Abdo, A. A., et al. 2010a, The Spectral Energy Distribution of Fermi bright blazars, Astrophys. J., 716, 30. arXiv:0912.2040

—. 2010b, Fermi Large Area Telescope First Source Catalog, Astrophys. J. Suppl., 188, 405. arXiv:1002.2280

—. 2011, Fermi large area telescope observations of Markarian 421: The missing piece of its spectral energy distribution, Astrophys. J., 736, 131. arXiv:1106.1348

Abreu, P., et al. 2012, Constraints on the origin of cosmic rays above 10^{18} eV from large scale anisotropy searches in data of the Pierre Auger Observatory, Astrophys. J., 762, L13. arXiv:1212.3083

Ackermann, M., et al. 2015a, The Third Catalog of Active Galactic Nuclei Detected by the Fermi Large Area Telescope, Astrophys. J., 810, 14. arXiv:1501.06054

—. 2015b, The spectrum of isotropic diffuse gamma-ray emission between 100 MeV and 820 GeV, Astrophys. J., 799, 86. arXiv:1410.3696

Aharonian, F. 2009, Simultaneous multiwavelength observations of the second exceptional gamma-ray flare of PKS 2155-304 in July 2006, Astron. Astrophys., 502, 749. arXiv:0906.2002

Aharonian, F., & Taylor, A. M. 2010, Limitations on the Photo-disintegration Process as a Source of VHE Photons, Astropart. Phys., 34, 258. arXiv:1005.3230

Aharonian, F. A. 2000, TeV gamma-rays from BL Lac objects due to synchrotron radiation of extremely high-energy protons, New Astron., 5, 377. arXiv:astro-ph/0003159

Ajello, M., et al. 2012, The Luminosity Function of Fermi-detected Flat-Spectrum Radio Quasars, Astrophys. J., 751, 108. arXiv:1110.3787

—. 2014, The Cosmic Evolution of Fermi BL Lacertae Objects, Astrophys. J., 780, 73. arXiv:1310.0006

Albert, A., et al. 2017, Search for High-energy Neutrinos from Binary Neutron Star Merger GW170817 with ANTARES, IceCube, and the Pierre Auger Observatory, Astrophys. J., 850, L35. arXiv:1710.05839

Alexander, K. D., et al. 2017, 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, L21. arXiv:1710.05457

Allard, D. 2012, Extragalactic propagation of ultrahigh energy cosmic-rays, Astropart. Phys., 39-40, 33. arXiv:1111.3290

Aloisio, R., Berezinsky, V., & Blasi, P. 2014, Ultra high energy cosmic rays: implications of Auger data for source spectra and chemical composition, JCAP, 1410, 020. arXiv:1312.7459

Aloisio, R., Berezinsky, V., Blasi, P., et al. 2007, A dip in the UHECR spectrum and the transition from galactic to extragalactic cosmic rays, Astropart. Phys., 27, 76. arXiv:astro-ph/0608219

Aloisio, R., Berezinsky, V., & Gazizov, A. 2012, Transition from galactic to extragalactic cosmic rays, Astropart. Phys., 39-40, 129. arXiv:1211.0494

Anchordoqui, L. A., Hooper, D., Sarkar, S., & Taylor, A. M. 2008, High-energy neutrinos from astrophysical accelerators of cosmic ray nuclei, Astropart. Phys., 29, 1. arXiv:astro-ph/0703001

Ansoldi, S., et al. 2018, The blazar TXS 0506+056 associated with a high-energy neutrino: insights into extragalactic jets and cosmic ray acceleration, Astrophys. J. Lett. arXiv:1807.04300

Apel, W. D., et al. 2011, Kneelike structure in the spectrum of the heavy component of cosmic rays observed with KASCADE-Grande, Phys. Rev. Lett., 107, 171104. arXiv:1107.5885

Araudo, A. T., Bell, A. R., & Blundell, K. M. 2015, Particle acceleration and magnetic field amplification in the jets of 4C74.26, Astrophys. J., 806, 243. arXiv:1505.02210

Araudo, A. T., Bell, A. R., Crilly, A., & Blundell, K. M. 2016, Evidence that the maximum electron energy in hotspots of FR II galaxies is not determined by synchrotron cooling, Mon. Not. Roy. Astron. Soc., 460, 3554. arXiv:1605.05125

Arcavi, I., et al. 2017. arXiv:1711.02671

Baerwald, P., Bustamante, M., & Winter, W. 2013, UHECR escape mechanisms for protons and neutrons from GRBs, and the cosmic ray-neutrino connection, Astrophys. J., 768, 186. arXiv:1301.6163

Baerwald, P., Hümmer, S., & Winter, W. 2012, Systematics in the Interpretation of Aggregated Neutrino Flux Limits and Flavor Ratios from Gamma-Ray Bursts, Astropart. Phys., 35, 508. arXiv:1107.5583

Baixeras, C., et al. 2004, Commissioning and first tests of the MAGIC telescope, Nucl. Instrum. Meth., A518, 188, doi: 10.1016/j.nima.2003.10.057

Baring, M. G. 1997, in Very high-energy phenomena in the universe. Proceedings, 32nd Rencontres de Moriond, Les Arcs, France, January 18-25, 1997, 97–106

Becker, J. K. 2008, High-energy neutrinos in the context of multimessenger physics, Phys. Rept., 458, 173. arXiv:0710.1557

Bell, A., Schure, K., Reville, B., & Giacinti, G. 2013, Cosmic ray acceleration and escape from supernova remnants, Mon. Not. Roy. Astron. Soc., 431, 415. arXiv:1301.7264

Bell, A. R. 2004, Turbulent amplification of magnetic field and diffusive shock acceleration of cosmic rays, "Mon. Not. Roy. Astron. Soc.", 353, 550, doi: 10.1111/j.1365-2966.2004.08097.x

Bell, A. R., Araudo, A. T., Matthews, J. H., & Blundell, K. M. 2018, Cosmic Ray Acceleration by Relativistic Shocks: Limits and Estimates, Mon. Not. Roy. Astron. Soc., 473, 2364. arXiv:1709.07793

Berezinsky, V., Gazizov, A. Z., & Grigorieva, S. I. 2006, On astrophysical solution to ultrahigh-energy cosmic rays, Phys. Rev., D74, 043005. arXiv:hep-ph/0204357

Biehl, D., Boncioli, D., Fedynitch, A., & Winter, W. 2017, Cosmic-Ray and Neutrino Emission from Gamma-Ray Bursts with a Nuclear Cascade. arXiv:1705.08909

Blandford, R., & Eichler, D. 1987, Particle Acceleration at Astrophysical Shocks: A Theory of Cosmic Ray Origin, Phys. Rept., 154, 1, doi: 10.1016/0370-1573(87)90134-7

Blandford, R. D., & Levinson, A. 1995, Pair cascades in extragalactic jets. 1: Gamma rays, Astrophys. J., 441, 79, doi: 10.1086/175338

Blandford, R. D., & Znajek, R. L. 1977, Electromagnetic extractions of energy from Kerr black holes, Mon. Not. Roy. Astron. Soc., 179, 433, doi: 10.1093/mnras/179.3.433

Bloom, S. D., & Marscher, A. P. 1996, An Analysis of the Synchrotron Self-Compton Model for the Multi–Wave Band Spectra of Blazars, Astrophys. J., 461, 657, doi: 10.1086/177092

Blumenthal, G. 1970, Energy loss of high-energy cosmic rays in pair-producing collisions with ambient photons, Phys. Rev., D 1, 1596, doi: 10.1103/PhysRevD.1.1596

Blumenthal, G. R., & Gould, R. J. 1970, Bremsstrahlung, synchrotron radiation, and compton scattering of high-energy electrons traversing dilute gases, Rev. Mod. Phys., 42, 237, doi: 10.1103/RevModPhys.42.237

Boettcher, M., & Dermer, C. D. 2010, Timing Signatures of the Internal-Shock Model for Blazars, Astrophys. J., 711, 445. arXiv:1001.1606

Boettcher, M., Reimer, A., Sweeney, K., & Prakash, A. 2013, Leptonic and Hadronic Modeling of Fermi-Detected Blazars, Astrophys. J., 768, 54. arXiv:1304.0605

Boncioli, D., Fedynitch, A., & Winter, W. 2017, Nuclear Physics Meets the Sources of the Ultra-High Energy Cosmic Rays, Scientific Reports, 7, 4882. arXiv:1607.07989

Bonning, E. W., Cheng, L., Shields, G. A., Salviander, S., & Gebhardt, K. 2007, Accretion Disk Temperatures and Continuum Colors in QSOs, Astrophys. J., 659, 211. arXiv:astro-ph/0611263

Burbidge, G. R. 1959, Estimates of the total energy in particles and magnetic field in the non-thermal radio source, Astrophys. J., 129, doi: 10.1086/146680

Catanese, M., et al. 1997, Multiwavelength observations of a flare from Markarian 501, Astrophys. J., 487, L143. arXiv:astro-ph/9707179

Cerruti, M., Zech, A., Boisson, C., et al. 2018a, in Week of French Astrophysics : Semaine de l'Astrophysique Francaise (SF2A 2018) Bordeaux, France, July 3-6, 2018

Cerruti, M., Zech, A., Boisson, C., et al. 2018b, Lepto-hadronic single-zone models for the electromagnetic and neutrino emission of TXS 0506+056. arXiv:1807.04335

Cerruti, M., Zech, A., Boisson, C., & Inoue, S. 2015, A hadronic origin for ultra-high-frequency-peaked BL Lac objects, MNRAS, 448, 910. arXiv:1411.5968

Chen, X., Pohl, M., Böttcher, M., & Gao, S. 2016, Particle diffusion and localized acceleration in inhomogeneous AGN jets – II. Stochastic variation, Mon. Not. Roy. Astron. Soc., 458, 3260. arXiv:1603.00900

Chiaberge, M., Capetti, A., & Celotti, A. 1999, The hst view of fr I radio galaxies: evidence for non-thermal nuclear sources, Astron. Astrophys., 349, 77. arXiv:astro-ph/9907064

Cleary, K., Lawrence, C. R., Marshall, J. A., Hao, L., & Meier, D. 2007, Spitzer Observations of 3C Quasars and Radio Galaxies: Mid-Infrared Properties of Powerful Radio Sources, Astrophys. J., 660, 117. arXiv:astro-ph/0612702

Coulter, D. A., et al. 2017, Swope Supernova Survey 2017a (SSS17a), the Optical Counterpart to a Gravitational Wave Source, Science. arXiv:1710.05452

Crank, J., & Nicolson, P. 1996, A practical method for numerical evaluation of solutions of partial differential equations of the heat-conduction type [reprint of MR0019410 (8,409b)], 6, 207, doi: https://doi.org/10.1007/BF02127704

D'Avanzo, P., et al. 2018, The evolution of the X-ray afterglow emission of GW 170817 / GRB 170817A in XMM-Newton observations, Astron. Astrophys., 613, L1. arXiv:1801.06164

de Gouveia Dal Pino, E. M., & Kowal, G. 2013, Particle Acceleration by Magnetic Reconnection. arXiv:1302.4374

De Marco, D., Blasi, P., & Olinto, A. V. 2003, On the Statistical significance of the GZK feature in the spectrum of ultrahigh-energy cosmic rays, Astropart. Phys., 20, 53. arXiv:astro-ph/0301497

De Ridder, S., Dvorak, E., & Gaisser, T. K. 2018, Sensitivity of IceCube Cosmic-Ray measurements to the hadronic interaction models, PoS, ICRC2017, 319, doi: 10.22323/1.301.0319

Dembinski, H. P., et al. 2019, in Ultra High Energy Cosmic Rays (UHECR 2018) Paris, France, October 8-12, 2018

Dermer, C., Ramirez-Ruiz, E., & Le, T. 2007, Correlation of Photon and Neutrino Fluxes in Blazars and Gamma Ray Bursts, Astrophys. J., 664, L67. arXiv:astro-ph/0703219

Dermer, C. D., & Schlickeiser, R. 1993, Model for the high-energy emission from blazars, Astrophys. J., 416, 458, doi: 10.1086/173251

Di Mauro, M., Cuoco, A., Donato, F., & Siegal-Gaskins, J. M. 2014, Fermi-LAT /*gamma*-ray anisotropy and intensity explained by unresolved Radio-Loud Active Galactic Nuclei, JCAP, 1411, 021. arXiv:1407.3275

Dolag, K., Grasso, D., Springel, V., & Tkachev, I. 2004, Mapping deflections of ultrahigh energy cosmic rays in constrained simulations of extragalactic magnetic fields, JETP Lett., 79, 583. arXiv:astro-ph/0310902

Drury, L. O. 1983, An introduction to the theory of diffusive shock acceleration of energetic particles in tenuous plasmas, Rept. Prog. Phys., 46, 973, doi: 10.1088/0034-4885/46/8/002

Elvis, M., Wilkes, B. J., McDowell, J. C., et al. 1994, Atlas of quasar energy distributions, Astrophys. J. Suppl., 95, 1, doi: 10.1086/192093

Essey, W., Kalashev, O. E., Kusenko, A., & Beacom, J. F. 2010, Secondary photons and neutrinos from cosmic rays produced by distant blazars, Phys. Rev. Lett., 104, 141102. arXiv:0912.3976

Fabian, A. C. 2012, Observational Evidence of AGN Feedback, Ann. Rev. Astron. Astrophys., 50, 455. arXiv:1204.4114

Fang, K., & Metzger, B. D. 2017, High-Energy Neutrinos from Millisecond Magnetars formed from the Merger of Binary Neutron Stars, Astrophys. J., 849, 153. arXiv:1707.04263

Farrar, G. R., & Sutherland, M. S. 2017, Deflections of UHECRs in the Galactic magnetic field. arXiv:1711.02730

Fermi, E. 1949, On the Origin of the Cosmic Radiation, Phys. Rev., 75, 1169, doi: 10.1103/PhysRev.75.1169

Fossati, G., Maraschi, L., Celotti, A., Comastri, A., & Ghisellini, G. 1998, A Unifying view of the spectral energy distributions of blazars, Mon. Not. Roy. Astron. Soc., 299, 433. arXiv:astro-ph/9804103

Fransson, C., & Bjornsson, C.-I. 1998, Radio emission and particle acceleration in sn 1993j, Astrophys. J., 509, 861. arXiv:astro-ph/9807030

Fransson, C., & Björnsson, C.-I. 2005, Modeling the radio and x-ray emission of SN 1993J and SN 2002ap, Springer Proc. Phys., 99, 59. arXiv:astro-ph/0404267

Gaisser, T. K., Stanev, T., & Tilav, S. 2013, Cosmic Ray Energy Spectrum from Measurements of Air Showers, Front. Phys.(Beijing), 8, 748. arXiv:1303.3565

Gallant, Y. A., & Achterberg, A. 1999, Ultra-high-energy cosmic ray acceleration by relativistic blast waves, Mon. Not. Roy. Astron. Soc., 305, 6. arXiv:astro-ph/9812316

Gao, S., Fedynitch, A., Winter, W., & Pohl, M. 2018, Interpretation of the coincident observation of a high energy neutrino and a bright flare, Nature Astronomy. arXiv:1807.04275

Gao, S., Pohl, M., & Winter, W. 2017, On the direct correlation between gamma-rays and PeV neutrinos from blazars, Astrophys. J., 843, 109. arXiv:1610.05306

Garrappa, S. 2018, in 8th International Fermi Symposium. https://fermi.gsfc.nasa.gov/science/mtgs/symposia/2018/program/tue/Simone_Garrappa.pdf

Gelmini, G. B., Kalashev, O., & Semikoz, D. V. 2012, Gamma-Ray Constraints on Maximum Cosmogenic Neutrino Fluxes and UHECR Source Evolution Models, JCAP, 1201, 044. arXiv:1107.1672

Georganopoulos, M., & Kazanas, D. 2003, Decelerating flows in TeV blazars: a resolution to the bl lac - fr I unification problem, Astrophys. J., 594, L27. arXiv:astro-ph/0307404

Ghisellini, G. 2010, The jet/disk connection in blazars, AIP Conf. Proc., 1242, 43. arXiv:1002.4619

—. 2016, The blazar sequence 2.0, Galaxies, 4, 36. arXiv:1609.08606

Ghisellini, G., & Madau, P. 1996, On the origin of the gamma-ray emission in blazars, Mon. Not. Roy. Astron. Soc., 280, 67, doi: 10.1093/mnras/280.1.67

Ghisellini, G., Righi, C., Costamante, L., & Tavecchio, F. 2017, The Fermi blazar sequence, Mon. Not. Roy. Astron. Soc., 469, 255. arXiv:1702.02571

Ghisellini, G., & Tavecchio, F. 2008, The blazar sequence: a new perspective, Mon. Not. Roy. Astron. Soc., 387, 1669. arXiv:0802.1918

Ghisellini, G., Tavecchio, F., & Chiaberge, M. 2005, Structured jets in TeV BL Lac objects and radiogalaxies: Implications on the observed properties, Astron. Astrophys., 432, 401. arXiv:astro-ph/0406093

Ghisellini, G., Tavecchio, F., Foschini, L., & Ghirlanda, G. 2011, The transition between BL Lac objects and Flat Spectrum Radio Quasars, Mon. Not. Roy. Astron. Soc., 414, 2674. arXiv:1012.0308

Ghisellini, G., Tavecchio, F., Foschini, L., et al. 2010, General physical properties of bright Fermi blazars, Mon. Not. Roy. Astron. Soc., 402, 497. arXiv:0909.0932

Giacinti, G., Kachelriess, M., Semikoz, D. V., & Sigl, G. 2012, Cosmic Ray Anisotropy as Signature for the Transition from Galactic to Extragalactic Cosmic Rays, JCAP, 1207, 031. arXiv:1112.5599

Giannios, D., Uzdensky, D. A., & Begelman, M. C. 2009, Fast TeV variability in blazars: jet in a jet, Mon. Not. Roy. Astron. Soc., 395, 29. arXiv:0901.1877

Giommi, P., Menna, M. T., & Padovani, P. 1999, The sedentary multi-frequency survey. I. statistical identification and cosmological properties of hbl bl lacs, Mon. Not. Roy. Astron. Soc., 310, 465. arXiv:astro-ph/9907014

Globus, N., Allard, D., Mochkovitch, R., & Parizot, E. 2015, UHECR acceleration at GRB internal shocks, Mon. Not. Roy. Astron. Soc., 451, 751. arXiv:1409.1271

Gokus, A., Richter, S., Spanier, F., et al. 2018, in IBWS Proceedings

Goldberg, D. E. 1989, Genetic Algorithms in Search, Optimization, and Machine Learning (New York: Addison-Wesley)

Goldstein, A., et al. 2017, An Ordinary Short Gamma-Ray Burst with Extraordinary Implications: Fermi-GBM Detection of GRB 170817A, Astrophys. J., 848, L14. arXiv:1710.05446

Gould, R. J., & Schreder, G. P. 1967, Pair Production in Photon-Photon Collisions, Phys. Rev., 155, 1404, doi: 10.1103/PhysRev.155.1404

Greene, J. E., & Ho, L. C. 2005, Estimating black hole masses in active galaxies using the H-alpha emission line, Astrophys. J., 630, 122. arXiv:astro-ph/0508335

Greisen, K. 1966, End to the cosmic ray spectrum?, Phys. Rev. Lett., 16, 748, doi: 10.1103/PhysRevLett.16.748

Haack, C., & Wiebusch, C. 2018, A measurement of the diffuse astrophysical muon neutrino flux using eight years of IceCube data., PoS, ICRC2017, 1005, doi: 10.22323/1.301.1005

Hallinan, G., et al. 2017, A Radio Counterpart to a Neutron Star Merger, Science, 358, 1579. arXiv:1710.05435

Halzen, F., & Kheirandish, A. 2016, High Energy Neutrinos from Recent Blazar Flares, Astrophys. J., 831, 12. arXiv:1605.06119

Heinze, J., Boncioli, D., Bustamante, M., & Winter, W. 2016, Cosmogenic Neutrinos Challenge the Cosmic Ray Proton Dip Model, Astrophys. J., 825, 122. arXiv:1512.05988

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

—. 2004, Where do 10**19-eV cosmic rays come from?, Nucl. Phys. Proc. Suppl., 136, 139, doi: 10.1016/j.nuclphysbps.2004.10.004

Hinton, J. A. 2004, The Status of the H.E.S.S. project, New Astron. Rev., 48, 331. arXiv:astro-ph/0403052

Hjorth, J., Levan, A. J., Tanvir, N. R., et al. 2017, The Distance to NGC 4993: The Host Galaxy of the Gravitational-wave Event GW170817, Astrophys. J., 848, L31. arXiv:1710.05856

Holler, M., et al. 2016, Observations of the Crab Nebula with H.E.S.S. Phase II, PoS, ICRC2015, 847. arXiv:1509.02902

Hooper, D., Linden, T., & Vieregg, A. 2018, Active Galactic Nuclei and the Origin of IceCube's Diffuse Neutrino Flux. arXiv:1810.02823

Hummer, S., Baerwald, P., & Winter, W. 2012, Neutrino Emission from Gamma-Ray Burst Fireballs, Revised, Phys. Rev. Lett., 108, 231101. arXiv:1112.1076

Hümmer, S., Rüger, M., Spanier, F., & Winter, W. 2010, Simplified models for photohadronic interactions in cosmic accelerators, Astrophys. J., 721, 630. arXiv:1002.1310

Inoue, Y., & Tanaka, Y. T. 2016, Baryon Loading Efficiency and Particle Acceleration Efficiency of Relativistic Jets: Cases For Low Luminosity BL Lacs, Astrophys. J., 828, 13. arXiv:1603.07623

Ishihara, A. 2016, A search for extremely high energy neutrinos in 6 years (2008-2014) of IceCube data, PoS, ICRC2015, 1064, doi: 10.22323/1.236.1064

Jones, F. C. 1994, A theoretical review of diffusive shock acceleration, "Astrophys. J. Suppl.", 90, 561, doi: 10.1086/191875

Jones, F. C., & Ellison, D. C. 1991, The plasma physics of shock acceleration, Space Sci. Rev., 58, 259, doi: 10.1007/BF01206003

Kadler, M., et al. 2016, Coincidence of a high-fluence blazar outburst with a PeV-energy neutrino event, Nature Phys., 12, 807. arXiv:1602.02012

Kampert, K.-H., Kulbartz, J., Maccione, L., et al. 2013, CRPropa 2.0 – a Public Framework for Propagating High Energy Nuclei, Secondary Gamma Rays and Neutrinos, Astropart. Phys., 42, 41. arXiv:1206.3132

Kaspi, S. 2001, Reverberation measurements of quasars and the size-mass-luminosity relations in active galactic nuclei, ASP Conf. Ser., 224, 347. arXiv:astro-ph/0009488

Kaspi, S., Maoz, D., Netzer, H., et al. 2005, The Relationship between luminosity and broad-line region size in active galactic nuclei, Astrophys. J., 629, 61. arXiv:astro-ph/0504484

Kawai, H., et al. 2008, Telescope array experiment, Nucl. Phys. Proc. Suppl., 175-176, 221, doi: 10.1016/j.nuclphysbps.2007.11.002

Keivani, A., et al. 2018, A Multimessenger Picture of the Flaring Blazar TXS 0506+056: implications for High-Energy Neutrino Emission and Cosmic Ray Acceleration, Astrophys. J., 864, 84. arXiv:1807.04537

Kimura, S. S., Murase, K., & Mészáros, P. 2018, Super-Knee Cosmic Rays from Galactic Neutron Star Merger Remnants, Astrophys. J., 866, 51. arXiv:1807.03290

Kimura, S. S., Murase, K., Mészáros, P., & Kiuchi, K. 2017, High-Energy Neutrino Emission from Short Gamma-Ray Bursts: Prospects for Coincident Detection with Gravitational Waves, Astrophys. J., 848, L4. arXiv:1708.07075

Konigl, A. 1981, Relativistic jets as X-ray and gamma-ray sources, Astrophys. J., 243, 700, doi: 10.1086/158638

Koning, A. J., Hilaire, S., & Duijvestijn, M. C. 2007, in Proceedings, International Conference on Nuclear Data for Science and Tecnology, 211–214

Kovalev, Y. Y., Lister, M. L., Homan, D. C., & Kellermann, K. I. 2007, The Inner Jet of the Radio Galaxy M87, Astrophys. J. Lett. arXiv:0708.2695

Krimm, H. A., et al. 2013, The Swift/BAT Hard X-ray Transient Monitor, Astrophys. J. Suppl., 209, 14. arXiv:1309.0755

Linsley, J. 1963, Evidence for a Primary Cosmic-Ray Particle with Energy 10^{20} eV, Physical Review Letters, 10, 146, doi: 10.1103/PhysRevLett.10.146

Lipunov, V. M., et al. 2017, MASTER Optical Detection of the First LIGO/Virgo Neutron Star Binary Merger GW170817, Astrophys. J., 850, L1. arXiv:1710.05461

Liu, H. T., & Bai, J. M. 2006, Absorption of 10-GeV to 200-GeV gamma rays by radiation from broad-line regions in blazars, Astrophys. J., 653, 1089. arXiv:0807.3135

Liu, R.-Y., Wang, K., Xue, R., et al. 2018, A hadronuclear interpretation of a high-energy neutrino event coincident with a blazar flare. arXiv:1807.05113

Loeb, A., & Waxman, E. 2006, The Cumulative background of high energy neutrinos from starburst galaxies, JCAP, 0605, 003. arXiv:astro-ph/0601695

Longair, M. S. 1994, High-energy astrophysics. Vol. 2: Stars, the galaxy and the interstellar medium

Lyman, J. D., et al. 2018, The optical afterglow of the short gamma-ray burst associated with GW170817, Nat. Astron., 2, 751. arXiv:1801.02669

Mannheim, K. 1993, The Proton blazar, Astron. Astrophys., 269, 67. arXiv:astro-ph/9302006

Maraschi, L., Ghisellini, G., & Celotti, A. 1992, A jet model for the gamma-ray emitting blazar 3C 279, Astrophys. J., 397, L5, doi: 10.1086/186531

Maraschi, L., & Tavecchio, F. 2003, The jet-disk connection and blazar unification, Astrophys. J., 593, 667. arXiv:astro-ph/0205252

Marchã, M. J. M., Browne, I. W. A., Impey, C. D., & Smith, P. S. 1996, Optical spectroscopy and polarization of a new sample of optically bright flat radio spectrum sources, MNRAS, 281, 425, doi: 10.1093/mnras/281.2.425

Margutti, R., et al. 2017, 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, L20. arXiv:1710.05431

—. 2018, The Binary Neutron Star Event LIGO/Virgo GW170817 160 Days after Merger: Synchrotron Emission across the Electromagnetic Spectrum, Astrophys. J., 856, L18. arXiv:1801.03531

Meszaros, P., & Waxman, E. 2001, TeV neutrinos from successful and choked gamma-ray bursts, Phys. Rev. Lett., 87, 171102. arXiv:astro-ph/0103275

Miller, B. L., Miller, B. L., Goldberg, D. E., & Goldberg, D. E. 1995, Genetic Algorithms, Tournament Selection, and the Effects of Noise, Complex Systems, 9, 193

Mooley, K. P., et al. 2018, A mildly relativistic wide-angle outflow in the neutron star merger GW170817, Nature, 554, 207. arXiv:1711.11573

Morejon, L. 2019, AstroPhoMes: Photomeson models for UHECR interactions with photons, doi: 10.5281/zenodo.2600178. https://doi.org/10.5281/zenodo.2600177

Mücke, A., Engel, R., Rachen, J. P., Protheroe, R. J., & Stanev, T. 2000, SOPHIA: Monte Carlo simulations of photohadronic processes in astrophysics, Comput. Phys. Commun., 124, 290. arXiv:astro-ph/9903478

Mucke, A., & Protheroe, R. J. 2001, A Proton synchrotron blazar model for flaring in Markarian 501, Astropart. Phys., 15, 121. arXiv:astro-ph/0004052

Muecke, A., Protheroe, R. J., Engel, R., Rachen, J. P., & Stanev, T. 2003, BL Lac Objects in the synchrotron proton blazar model, Astropart. Phys., 18, 593. arXiv:astro-ph/0206164

Murase, K. 2007, High energy neutrino early afterglows gamma-ray bursts revisited, Phys. Rev., D76, 123001. arXiv:0707.1140

Murase, K., & Beacom, J. F. 2010, Very-High-Energy Gamma-Ray Signal from Nuclear Photodisintegration as a Probe of Extragalactic Sources of Ultrahigh-Energy Nuclei, Phys. Rev., D82, 043008. arXiv:1002.3980

Murase, K., Guetta, D., & Ahlers, M. 2016, Hidden Cosmic-Ray Accelerators as an Origin of TeV-PeV Cosmic Neutrinos, Phys. Rev. Lett., 116, 071101. arXiv:1509.00805

Murase, K., Inoue, Y., & Dermer, C. D. 2014, Diffuse Neutrino Intensity from the Inner Jets of Active Galactic Nuclei: Impacts of External Photon Fields and the Blazar Sequence, Phys. Rev., D90, 023007. arXiv:1403.4089

Murase, K., & Ioka, K. 2013, TeV–PeV Neutrinos from Low-Power Gamma-Ray Burst Jets inside Stars, Phys. Rev. Lett., 111, 121102. arXiv:1306.2274

Murase, K., Oikonomou, F., & Petropoulou, M. 2018, Blazar Flares as an Origin of High-Energy Cosmic Neutrinos?, Astrophys. J., 865, 124. arXiv:1807.04748

Nalewajko, K. 2013, The brightest gamma-ray flares of blazars, Mon. Not. Roy. Astron. Soc., 430, 1324. arXiv:1211.0274

Neronov, A. Y., Semikoz, D. V., & Tkachev, I. I. 2009, Ultra-High Energy Cosmic Ray production in the polar cap regions of black hole magnetospheres, New J. Phys., 11, 065015. arXiv:0712.1737

Nynka, M., Ruan, J. J., & Haggard, D. 2018. arXiv:1805.04093

Paczynski, B., & Xu, G. H. 1994, Neutrino bursts from gamma-ray bursts, Astrophys. J., 427, 708, doi: 10.1086/174178

Padovani, P., Giommi, P., Resconi, E., et al. 2018, Dissecting the region around IceCube-170922A: the blazar TXS 0506+056 as the first cosmic neutrino source, Mon. Not. Roy. Astron. Soc., 480, 192. arXiv:1807.04461

Padovani, P., Oikonomou, F., Petropoulou, M., Giommi, P., & Resconi, E. 2019, TXS 0506+056, the first cosmic neutrino source, is not a BL Lac. arXiv:1901.06998

Padovani, P., Resconi, E., Giommi, P., Arsioli, B., & Chang, Y. L. 2016, Extreme blazars as counterparts of IceCube astrophysical neutrinos, Mon. Not. Roy. Astron. Soc., 457, 3582. arXiv:1601.06550

Palladino, A., Rodrigues, X., Gao, S., & Winter, W. 2019, Interpretation of the diffuse astrophysical neutrino flux in terms of the blazar sequence, Astrophys. J., 871, 41. arXiv:1806.04769

Palladino, A., Spurio, M., & Vissani, F. 2016, On the IceCube spectral anomaly, JCAP, 1612, 045. arXiv:1610.07015

Palladino, A., & Vissani, F. 2016, Extragalactic plus Galactic model for IceCube neutrino events, Astrophys. J., 826, 185. arXiv:1601.06678

—. 2017, Can BL Lacertae emission explain the neutrinos above 0.2 PeV?, Astron. Astrophys., 604, A18. arXiv:1702.08779

Palladino, A., & Winter, W. 2018, A multi-component model for observed astrophysical neutrinos, Astron. Astrophys., 615, A168. arXiv:1801.07277

Panessa, F., Baldi, R. D., Laor, A., et al. 2019, The Origin of Radio Emission from Radio-Quiet AGN. arXiv:1902.05917

Peters, B. 1961, Primary cosmic radiation and extensive air showers, Il Nuovo Cimento, 22, 800, doi: 10.1007/BF02783106

Piro, A. L., & Kollmeier, J. A. 2018, Evidence for Cocoon Emission from the Early Light Curve of SSS17a, Astrophys. J., 855, 103. arXiv:1710.05822

Protheroe, R. J. 1995, High-energy particles from active galactic nuclei, Nucl. Phys. Proc. Suppl., 43, 229, doi: 10.1016/0920-5632(95)00479-S

Reimer, A., Boettcher, M., & Buson, S. 2018, Cascading Constraints from Neutrino Emitting Blazars: The case of TXS 0506+056. arXiv:1812.05654

Resconi, E., Coenders, S., Padovani, P., Giommi, P., & Caccianiga, L. 2017, Connecting blazars with ultrahigh-energy cosmic rays and astrophysical neutrinos, Mon. Not. Roy. Astron. Soc., 468, 597. arXiv:1611.06022

Righi, C., Tavecchio, F., & Guetta, D. 2017, High-energy emitting BL Lacs and high-energy neutrinos - Prospects for the direct association with IceCube and KM3NeT, Astron. Astrophys., 598, A36. arXiv:1607.08061

Rodrigues, X., Biehl, D., Boncioli, D., & Taylor, A. M. 2019, Binary neutron star merger remnants as sources of cosmic rays below the “Ankle”, Astropart. Phys., 106, 10. arXiv:1806.01624

Rodrigues, X., Fedynitch, A., Gao, S., Boncioli, D., & Winter, W. 2018a, Neutrinos and Ultra-High-Energy Cosmic-Ray Nuclei from Blazars, Astrophys. J., 854, 54. arXiv:1711.02091

Rodrigues, X., Gao, S., Fedynitch, A., Palladino, A., & Winter, W. 2018b, Leptohadronic blazar models applied to the 2014–15 flare of TXS 0506+056, accepted for publication in Astrophys. J. Lett. arXiv:1812.05939

Ruan, J. J., Nynka, M., Haggard, D., Kalogera, V., & Evans, P. 2018, Brightening X-Ray Emission from GW170817/GRB 170817A: Further Evidence for an Outflow, Astrophys. J., 853, L4. arXiv:1712.02809

Sadowski, A., & Narayan, R. 2015, Powerful radiative jets in supercritical accretion discs around non-spinning black holes, Mon. Not. Roy. Astron. Soc., 453, 3213. arXiv:1503.00654

Sahakyan, N. 2018, Lepto-hadronic γ-ray and neutrino emission from the jet of TXS 0506+056, Astrophys. J., 866, 109. arXiv:1808.05651

Savchenko, V., et al. 2017, INTEGRAL Detection of the First Prompt Gamma-Ray Signal Coincident with the Gravitational-wave Event GW170817, Astrophys. J., 848, L15. arXiv:1710.05449

Sigl, G., Miniati, F., & Ensslin, T. A. 2004, Ultrahigh energy cosmic ray probes of large scale structure and magnetic fields, Phys. Rev., D70, 043007. arXiv:astro-ph/0401084

Sikora, M., Begelman, M. C., & Rees, M. J. 1994, Comptonization of diffuse ambient radiation by a relativistic jet: The source of gamma rays from blazars?, Astrophys. J., 421, 153, doi: 10.1086/173633

Sikora, M., Blazejowski, M., Moderski, R., & Madejski, G. M. 2002, On the nature of MeV blazars, Astrophys. J., 577, 78. arXiv:astro-ph/0205527

Sikora, M., Rutkowski, M., & Begelman, M. 2016, A spine–sheath model for strong-line blazars, Mon. Not. Roy. Astron. Soc., 457, 1352. arXiv:1511.08924

Sinnis, G., Smith, A., & McEnery, J. E. 2004, in On recent developments in theoretical and experimental general relativity, gravitation, and relativistic field theories. Proceedings, 10th Marcel Grossmann Meeting, MG10, Rio de Janeiro, Brazil, July 20-26, 2003. Pt. A-C, 1068–1088

Soares-Santos, M., et al. 2017, The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/Virgo GW170817. I. Discovery of the Optical Counterpart Using the Dark Energy Camera, Astrophys. J., 848, L16. arXiv:1710.05459

Spada, M., Ghisellini, G., Lazzati, D., & Celotti, A. 2001, Internal shocks in the jets of radio-loud quasars, Mon. Not. Roy. Astron. Soc., 325, 1559. arXiv:astro-ph/0103424

Stecker, F. W. 2013, PeV neutrinos observed by IceCube from cores of active galactic nuclei, Phys. Rev., D88, 047301. arXiv:1305.7404

Strotjohann, N. L., Kowalski, M., & Franckowiak, A. 2018, Eddington Bias for Cosmic Neutrino Sources. arXiv:1809.06865

Takami, H., Kyutoku, K., & Ioka, K. 2014, High-Energy Radiation from Remnants of Neutron Star Binary Mergers, Phys. Rev., D89, 063006. arXiv:1307.6805

Tamborra, I., Ando, S., & Murase, K. 2014, Star-forming galaxies as the origin of diffuse high-energy backgrounds: Gamma-ray and neutrino connections, and implications for starburst history, JCAP, 1409, 043. arXiv:1404.1189

Tatischeff, V. 2009, Radio emission and nonlinear diffusive shock acceleration of cosmic rays in the supernova SN 1993J, Astron. Astrophys., 499, 191. arXiv:0903.2944

Tavecchio, F., & Ghisellini, G. 2008, The spectrum of the Broad Line Region and the high-energy emission of powerful blazars, Mon. Not. Roy. Astron. Soc., 386, 945. arXiv:0802.0871

—. 2016, On the magnetization of BL Lac jets, Mon. Not. Roy. Astron. Soc., 456, 2374. arXiv:1509.08710

Tavecchio, F., Ghisellini, G., & Guetta, D. 2014, Structured jets in BL Lac objects: efficient PeV neutrino factories?, Astrophys. J., 793, L18. arXiv:1407.0907

Troja, E., et al. 2017, The X-ray counterpart to the gravitational wave event GW 170817, Nature, 551, 71. arXiv:1710.05433

Unger, M. 2018, Highlights from the Pierre Auger Observatory, PoS, ICRC2017, 1102. arXiv:1710.09478

Unger, M., Farrar, G. R., & Anchordoqui, L. A. 2015, Origin of the ankle in the ultrahigh energy cosmic ray spectrum, and of the extragalactic protons below it, Phys. Rev., D92, 123001. arXiv:1505.02153

Urry, C. M., & Padovani, P. 1995, Unified schemes for radio-loud active galactic nuclei, Publ. Astron. Soc. Pac., 107, 803. arXiv:astro-ph/9506063

Valenti, S., Sand, D. J., Yang, S., et al. 2017, The discovery of the electromagnetic counterpart of GW170817: kilonova AT 2017gfo/DLT17ck, Astrophys. J., 848, L24. arXiv:1710.05854

van Eerten, E. T. H., Ryan, G., Ricci, R., et al. 2018, A year in the life of GW170817: the rise and fall of a structured jet from a binary neutron star merger. arXiv:1808.06617

Vietri, M. 1998, Ultrahigh-energy neutrinos from gamma-ray bursts, Phys. Rev. Lett., 80, 3690. arXiv:astro-ph/9802241

Villar, V. A., et al. 2018, *SpitzerSpaceTelescope* Infrared Observations of the Binary Neutron Star Merger GW170817, Astrophys. J., 862, L11. arXiv:1805.08192

Villata, M., & Raiteri, C. M. 1999, Helical jets in blazars. I. The case of MKN 501, Astronomy & Astrophysics, 347, 30

Wagner, S. J., & Witzel, A. 1995, Intraday variability in quasars and bl lac objects, Ann. Rev. Astron. Astrophys., 33, 163, doi: 10.1146/annurev.aa.33.090195.001115

Wang, B., & Li, Z. 2016, Can FSRQs produce the IceCube detected diffuse neutrino emission?, Sci. China Phys. Mech. Astron., 59, 619502. arXiv:1505.04418

Waxman, E. 1995, Cosmological origin for cosmic rays above 10**19-eV, Astrophys. J., 452, L1. arXiv:astro-ph/9508037

Waxman, E., & Bahcall, J. N. 1997, High-energy neutrinos from cosmological gamma-ray burst fireballs, Phys. Rev. Lett., 78, 2292. arXiv:astro-ph/9701231

—. 1999, High-energy neutrinos from astrophysical sources: An Upper bound, Phys. Rev., D59, 023002. arXiv:hep-ph/9807282

Weekes, T. C., et al. 2002, VERITAS: The Very energetic radiation imaging telescope array system, Astropart. Phys., 17, 221. arXiv:astro-ph/0108478

Yu, X., Zhang, X., Zhang, H., et al. 2015, The jets-accretion relation, mass-luminosity relation in Fermi blazars, Astrophys. Space Sci., 357, 14. arXiv:1504.05635

Zatsepin, G. T., & Kuzmin, V. A. 1966, Upper limit of the spectrum of cosmic rays, JETP Lett., 4, 78

Zhang, B. T., & Li, Z. 2017, Constraints on cosmic ray loading and PeV neutrino production in blazars, JCAP, 1703, 024. arXiv:1607.02211

Zhang, B. T., Murase, K., Oikonomou, F., & Li, Z. 2017, High-energy cosmic ray nuclei from tidal disruption events: origin, survival, and implications. arXiv:1706.00391

Zhang, H., Fang, K., & Li, H. 2018, Variability and Optical Polarization Can Probe the Neutrino and Electromagnetic Emission Mechanisms of TXS 0506+056. arXiv:1807.11069

Appendix A

Model parameters

In this appendix we detail the parameter values for the sources considered in the population study in Chapter 4 (Palladino et al., 2019, Tab. A.1), and for the TXS models presented in Chapter 5 (Rodrigues et al., 2018b, Tab. A.2).

Model ID	Type	$N_{\rm sources}$	L_γ $\log_{10}$, erg/s	$B'_{\rm jet}$ mG	$r_{\rm BLR}$ pc	$r_{\rm DT}$ pc
11	HL-FSRQ	49	50.3	5000	2.54	64
10			49.6	2300	0.73	16
9 (B)			48.8	900	0.18	4.0
8	LBL/FSRQ	384	47.7	260	–	–
7			46.1	40	–	–
6	IBL/FSRQ	285	45.7	28	–	–
5			45.5	21	–	–
4 (A)			44.6	7.1	–	–
3	HBL	29	43.5	2.2	–	–
2			42.5	0.69	–	–
1			41.6	0.23	–	–

Table A.1: Parameters of the models shown in Chapter 3 for the blazar sequence (Figure 4.2). $N_{\rm sources}$ refers to the number of observed blazars with measured redshift in Fermi-3LAC catalog (Ackermann et al., 2015a). Note that among model ID 4 to 8, there are large overlaps between the FSRQ and BL Lac populations. From model ID 1-8, we ignored the BLR and dust torus radiation zones, since they are either negligible in all BL Lacs or the blob in the jet lies outside the regions in those FSRQs. The last columns show the results from the simulations (obtained neutrino-to-gamma-ray luminosity ratio for a CR loading $\xi_{\rm CR} = 10$ and considering diffusive CR escape).

Table A.2: Selected parameter sets for each model discussed in Chapter 5, along with the predicted number of neutrino events N_ν and the SED quality compared to data. The values for 1-zone (a) and (b) are given for two representative curves from Figure 5.1 (red curve from the left panel and green curve from the right panel). The physical luminosities $L^{\rm obs}_{e,\rm phys}$ and $L^{\rm obs}_{p,\rm phys}$ carried by electrons and protons are given by $L^{\rm obs}_{\rm phys} = L_{\rm iso}/\Gamma^2$, where $L_{\rm iso}$ is the isotropic-equivalent luminosity. The physical luminosities can be compared to the Eddington luminosity, which is 4×10^{46} erg/s for a black hole mass of 3×10^8 $M_\odot$ estimated by Padovani et al. (2019); note, however, that this value can be temporarily exceeded during flares. Table taken from Rodrigues et al. (2018b).

	Quality criteria		Parameters									
Model	N_ν	SED	B' [G]	$R'_{\rm blob}$ [cm]	$\Gamma_{\rm b}$	$\gamma^{\rm max}_{\rm p,\,obs}$	$L^{\rm obs}_{e,\rm phys}$ [erg/s]	$L^{\rm obs}_{p,\rm phys}$ [erg/s]	$f_{\rm esc}$	$T_{\rm disk}$ [K]	$L_{\rm disk}$ [erg/s]	$R_{\rm BLR}$ [cm]
1-zone (a)	1.8	Compatible	1.0	10^{17}	9.0	$10^{5.8}$	$10^{44.2}$	$10^{49.6}$	$10^{-2.5}$	–	–	–
1-zone (b)	13.2	Overshoot	0.001	10^{15}	7.0	$10^{5.7}$	$10^{44.8}$	$10^{50.7}$	$10^{-2.5}$	–	–	–
C. core (blob)	0.0	Compatible	0.01	$10^{18.7}$	10.0	–	$10^{44.8}$	–	$10^{-2.5}$	–	–	–
C. core (core)	1.9			10^{15}		$10^{6.1}$	$10^{43.7}$	$10^{49.5}$	$10^{-2.5}$	–	–	–
Ext. field (a)	4.9	Compatible	0.6	$10^{15.8}$	49.1	$10^{5.9}$	$10^{43.6}$	$10^{48.7}$	$10^{-4.8}$	$10^{5.7}$	$10^{44.7}$	$10^{17.8}$
Ext. field (b)	4.0	Cutoff, 10 GeV	0.9	$10^{16.3}$	48.0	$10^{6.3}$	$10^{42.9}$	$10^{48.4}$	$10^{-3.1}$	$10^{5.3}$	$10^{44.7}$	$10^{17.3}$

Appendix B

Nuclear cascade setup and numerical methods of NeuCosmA

In this appendix we outline the basic methods involved in the lepto-hadronic blazar model used in Chapters 3 and 4. It is based on the NeuCosmA code, which was applied past studies to GRBs (Boncioli et al., 2017; Biehl et al., 2017). As mentioned in the main text, in this approach only hadronic species are calculated explicitly, while the photon distribution, considered constant throughout the simulation, is taken from observations and is assumed to be maintained by a population of co-accelerated electrons. We have adapted some aspects of the method in order to model blazar flares. For instance, instead of an approach where the particle densities are allowed to reach a steady state, we integrate the time-dependent escaping spectra of the hadronic species, as well as neutrinos (Rodrigues et al., 2018a), until the injection of the primary hadrons stops and eventually all the densities in the source drop exponentially to zero. All these different aspects are discussed below.

B.1 Nuclear system setup

Starting with the heaviest possible injection isotope (iron-56 in the case illustrated in Figure B.1), the isotopes created through photo-hadronic interactions and decay are determined, by following all possible paths through a recursive algorithm. The most stable isotopes (closer to the "valley" in the middle of the chart, as discussed in Section 3.1) will be dominantly populated. The isotopes marked in blue and green in Figure B.1 were included explicitly in the coupled differential equation system (discussed in the next section). The remaining isotopes, not marked in green or blue, will not be created in the source, and are therefore not included in the calculation.

As mentioned in the main text, the radiation processes included for the nuclei are $\beta^{\pm}$ decays, spontaneous emission of nucleons and light nuclei, and photo-disintegration (modeled with TALYS (Koning et al., 2007) for $A \geq 12$ and CRPropa2 (Kampert et al., 2013) for $A < 12$). In addition, all the other photo-hadronic and electromagnetic processes considered for protons are also included for nuclei, namely, Bethe-Heitler photo-pair production (Blumenthal, 1970), syn-

chrotron cooling, escape (which depends on the CR escape assumption, *cf.* Section 2.1.3), and photo-meson production. Like for protons, photo-meson production is based on the SOPHIA code (Muecke et al., 2003); the extension to nuclei is based on an extension of the model by Hümmer et al. (2010), assuming the interaction cross section scales with the mass of the isotope, $\sigma_{p\gamma} \propto A$.

The isotopes marked by crosses are beta emitters with a lifetime lower than 1 second, and were determined to be relevant for neutrino production through beta decay. This is because an UHECR with lifetime of $\tau_0 = 1$ s moving with a Lorentz factor of $\gamma' = 10^6$ will travel a distance in the shock rest frame of $\gamma' c\tau_0 = 3 \times 10^{16}$ cm before undergoing beta-decay, which is the size of the radiation zone considered in Chapters 3 and 4. Therefore, any species that decay slower than this threshold will not contribute considerably to neutrino production in the jet through beta decay, which is why this threshold was chosen. Note that only green and blue isotopes are included in the calculation, which means that neutrino emission from beta decay is only relevant from the green, crossed isotopes.

All spontaneous nucleon or α-particle emitters faster than a certain threshold (boxes with dots in Figure B.1 are integrated out in the beginning, which means that they are not explicitly considered in the PDE system but are instead replaced by their daughters. This selection is made conservatively, *i.e.* using a low threshold value of $\tau_0 < 10^{-10}$ s, where τ_0 is the decay time of the nucleus in its own rest frame. That way, we guarantee that no other processes are being neglected that could otherwise compete with the spontaneous emission.

B.2 Numerical solution

The NeuCosmA code integrates the system of partial differential equations (PDEs) describing the evolution of the density spectrum of each hadronic species in the source (protons, neutrons and nuclei). In what follows, the particle density spectrum of species i given in the shock rest frame is represented by the notation $N'_i \equiv (1/E')dN'_i/dE'$ [GeV^{-1}cm^{-3}]. The generic form of the PDE describing the evolution of species i is

$$\frac{\partial N'_i}{\partial t'} = \frac{\partial}{\partial E'}\left(-b'(E')N'_i(E')\right) - \frac{N'_i(E')}{t'_{\rm esc}(E')} + \tilde{Q}'(E')\,, \tag{B.1}$$

where $b'(E) = E' t'^{-1}_{\rm loss}(E')$, with $t'^{-1}_{\rm loss}(E') = E'^{-1}|dE'/dt|$ [s^{-1}] the total energy loss rate, and $t'_{\rm esc}$ is the escape timescale. The first term contains all the continuous energy losses of the particles; the escape term contains not only the physical escape from the region (*cf.* Sections 2.1.3 and 3.4), but also a "sink term" caused by interactions that change particle i into a different species. That is the case, for example, for photo-disintegration or $\beta^{\pm}$ decay. The last term is the injection term

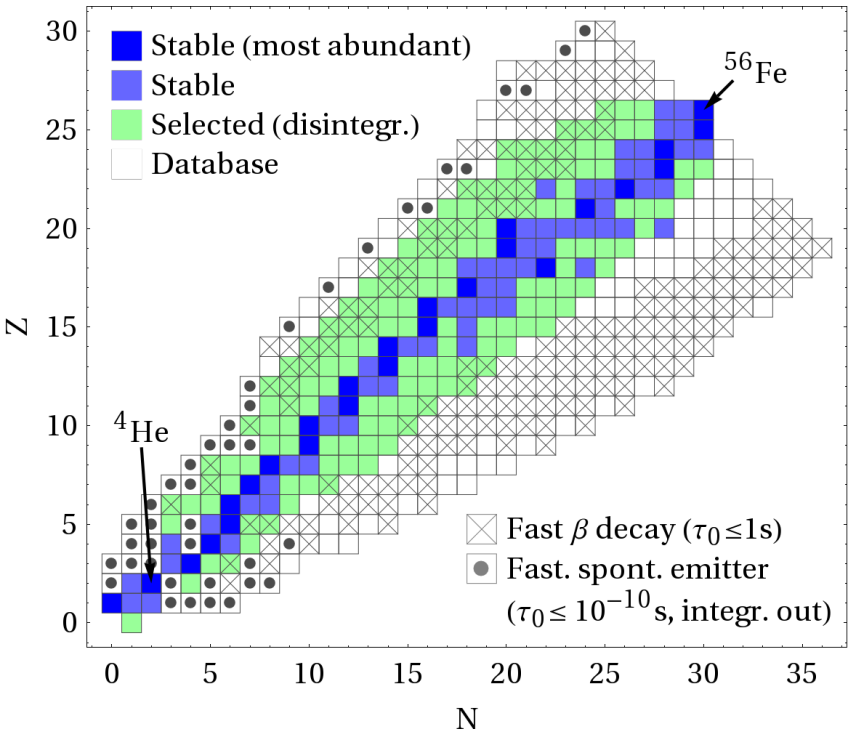

Figure B.1: Nuclear isotopes considered in this study as a function of neutron number N and proton number Z ("database", white boxes). The fast spontaneous emitters marked by dots are integrated out immediately, while the colored isotopes are selected by a recursive algorithm following the leading disintegration and decay paths. The blue isotopes are stable, while the green ones are typically unstable $\beta^{\pm}$ emitters. Fast $\beta^{\pm}$ emitters (with lifetimes $\tau_0 < 1$ s, such that they could potentially contribute to neutrino production in the jet) are marked with crosses. Figure taken from Rodrigues et al. (2018a).

(the notation $Q' = dN'_{\rm inj}/dt'$ [GeV^{-1}cm^{-3}] is used to represent the energy density of particles injected per unit time),

$$\tilde{Q}'(E') = Q'_i(E') + \sum_j Q'_{j\to i}(E')\,. \tag{B.2}$$

where Q'_i describes the injection from an acceleration zone, which is the case only for the primary species injected in each simulation; the component $\sum_j Q'_{j\to i}$, describes the "injection" from interactions of particles of all other species j, that mutate them into species i.

The PDE system is then integrated numerically over a large number of small time steps. The NeuCosmA code solves this system using the Crank-Nicolson method (Crank & Nicolson, 1996), an efficient solver that is second-order in time. The method is implicit in time, which means that the variable t' in Eq. (B.1) does not appear explicitly; instead, the densities of the successive time steps are calculated in an iterative fashion. The actual variables that the system is solved for are $E'^2 N'_i(E')$ [GeV cm^{-3}].

Throughout this work we considered only pure injection compositions (protons, helium-4 and iron-56 in Chapter 3, and only protons in Chapter 4. This means that there is only one term Q'_iin the PDE system, belonging to the injected species. As mentioned in Section 3.2, a simple power law was considered as the injection spectrum for the primary CRs, with a super-exponential cut-off at maximum energy:

$$Q'_i \propto E'^{-2}_i \cdot \exp\left(-\left(E'_i/E'_{i,\rm max}\right)^2\right). \tag{B.3}$$

As explained in Section 3.2, the maximum energy $E'_{i,\rm max}$ is determined by equating the acceleration rate (Eq. 2.1) with the sum of synchrotron loss, photo-disintegration and photo-meson production rates, when the source is optically thick, or with the inverse light-crossing time of the blob, $c/R'_{\rm blob}$, when the source is optically thin.

B.3 Integration of escaping particles

In other previous works (Boncioli et al., 2017; Biehl et al., 2017) where the NeuCosmA code was applied to the study of GRBs, the system is first evolved to a steady state, and only then are the CR and neutrino fluxes calculated. On the contrary, the results discussed in Chapter 3 were obtained using a time-dependent computation for the fluence of CRs and neutrinos, *i.e.* the escaping fluxes are explicitly integrated in order to obtain the total fluxes emitted during the flare. The solver integrates the blob system up to a predetermined simulation duration $t'_{\rm simul} > t'_{\rm flare}$, while the injection of nuclei is switched off at $t'_{\rm flare}$ (and the electrons are assumed to maintain the non-thermal photon field). After the CR injection is switched off, the remaining

CRs in the system will cool or escape, and the particle densities will decrease exponentially. The simulation duration is determined so that the majority of the CRs have time to escape, and the remaining densities in the system are negligible. For advective escape this decay occurs relatively fast and requires a simulation duration of only a few $t'_{\rm flare}$, while for diffusive escape the simulation must be run for longer times (*cf.* Section 2.1.3).

At every time step of the integration of the PDE system, the spectrum of neutrinos and CRs that escape the jet during that time step is added to the total spectrum. Since we are interested in the total number of particles, the densities must be integrated over the volume of the blob, which corresponds simply to $V'_{\rm blob} N'_i$ [GeV^{-1}], since all the particles are considered to be distributed homogeneously. We then obtain:

$$V'_{\rm blob}\, N_i'^{\rm esc}(E') \;=\; \sum_{t'=0}^{t'_{\rm simul}} \frac{dN_i'^{\rm esc}(E',t')}{dt'}\,\Delta t' \;=\; \sum_{t'=0}^{t'_{\rm simul}} \frac{N'_i(E',t')\,\Delta t'}{t'_{\rm esc}(E')}, \tag{B.4}$$

where $\Delta t'$ is the integration time step (in units of physical time and in the blob frame), and $t'_{\rm esc}$ is the escape timescale. For neutrinos, the escaping spectrum from the jet is simply stored, since it will contribute to the total emitted spectrum. For CRs, the spectra are transformed into density spectra in the BLR, in order to compute the interaction sin that zone:

$$N_i^{\rm BLR}(E) \;=\; \frac{\Gamma_{\rm b}}{V_{\rm BLR}}\, V'_{\rm blob}\, N_i'^{\rm esc}(E'), \tag{B.5}$$

where $V_{\rm BLR}$ is the volume of the BLR zone. Note the Lorentz boost of the CR and neutrinos energies, implied by the change of variable from E' to $E = \Gamma E'$ and the inclusion of the bulk Lorentz factor of the jet $\Gamma_{\rm b}$ (following an argument similar to Eq. (2.18)). We now solve a new PDE system for the BLR, but starting with the CR densities ejected from the jet, $N_i^{\rm BLR}(E, t = t_{\rm flare}) > 0$. Since there is no further injection of primary CRs, all species have zero injection term $Q'_i(E^prime)$ (Eq. B.3), and the CR densities will decrease through interactions with the external photons. Escape into the dust torus zone occurs with the free-streaming timescale, $t_{\rm esc} = R_{\rm BLR}/c$, since no magnetic fields are considered present in the BLR. The dust torus region is treated in a similar manner, starting with the integrated CR spectra that escaped from the BLR.

Note that this approach is more expensive than a steady-state solution in terms of computing time, and in fact yields similar results for the jet emission (at least given the parameter values explored in this study); however, it is essential for the modeling of the CR transport between the jet, BLR and DT zones in high-luminosity FSRQs (*cf.* Section 3.5).

In Figure B.2 we show the total (energy-integrated) luminosity of the CRs escaping from the three zones of a high-luminosity FSRQ, as a function of time, for the case of diffusive escape

(the corresponding spectra are shown in Figure 3.7). Note that the timescales of the successive zones are orders of magnitude larger (the t-axis is logarithmic). In this example we consider iron-56 injection. As we can see, the luminosity of iron-56 emitted from the jet (left part of the plot) grows slower than secondary species, which is due to efficient photo-hadronic processes. At $t = t_{\rm flare} = 10^5$ s, the injection of iron halts and its emission levels start decreasing exponentially, while other CRs and neutrinos continue being emitted abundantly form the jet up to $t \sim 10 t_{\rm flare}$, due to the continuing hadronic processes acting on the secondary nuclei. In the BLR and dust torus, the emitted luminosity is lower due to the long escape timescales. However, the time-integrated spectrum emitted from these zones can still be considerable (*cf.* Figure 3.9).

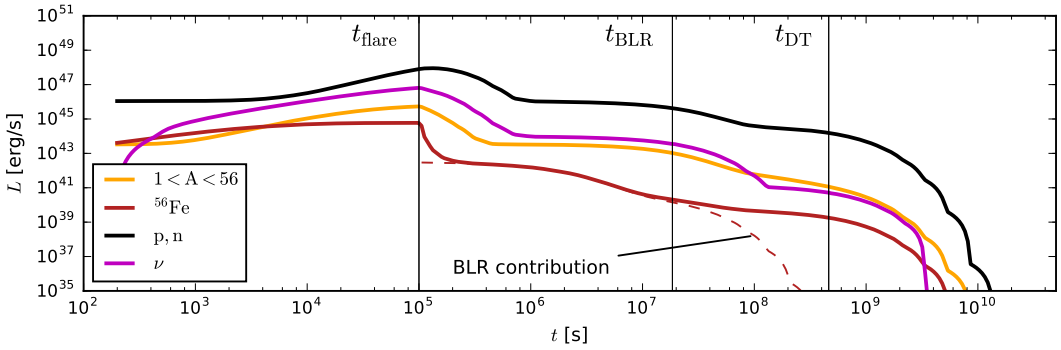

Figure B.2: Evolution of the CR and neutrino luminosities emitted from the successive zones of the high-luminosity FSRQ model (Section 3.5). We consider iron-56 injection in a blazar with $L_\gamma = 10^{48.8}$ erg s^{-1}, and a diffusive escape mechanism. Note that at $t_{\rm flare}$ the iron-56 injection stops, and the subsequent emission is due to the continuing interactions of secondary particles in the subsequent zones. Figure taken from Rodrigues et al. (2018a).

Appendix C

Genetic algorithm for TXS 0506+056 model optimization

Here we discuss the most important details of the algorithm used in Rodrigues et al. (2018b) for optimizing the parameters of the external field model of TXS 0506+056, in order to best describe the multi-wavelength data and the neutrino signal observed in 2014–15.

The method is based on a genetic algorithm, whose basic concepts were laid out by Goldberg (1989), for instance. The code was written in Python 2.7 using evolutionary functions from the DEAP module[1].

The idea of the algorithm is akin to the mechanisms of natural selection in biology: we start by creating a large, random population across a large portion of the parameter space, and let it develop through several generations, subject to selective evolution. The selection criteria must reflect the quality of each individual compared to our ideal, optimized version, so that the population evolves towards the location(s) in parameter space that yield optimal results.

In the problem at hand, each "individual" consists a set of seven parameters that define the physical environment of the blazar within the external field model, as discussed in Section 5.3. The evaluation of each individual is done by running the lepto-hadronic simulation with the respective parameter set to obtain the predicted SED and number of muon tracks in IceCube, and then comparing them with the multi-wavelength constraints and the IceCube signal event. Expanding on the analogy with biology, the parameter set can be regarded as the "genotype" of the individual, which we are looking to optimize, and the SED and neutrino spectrum the "phenotype", from which the "fitness" of the individual is computed.

Fitness evaluation– Our fitness function is defined as follows: first, the goodness of the emitted SED is evaluated against the multi-wavelength data constraints, based on a χ^2 criterion:

$$\chi^2_{\rm SED} = \sum_i \frac{(\mathcal{F}_i^{\rm model} - \mathcal{F}_i^{\rm obs})^2}{\sigma_i^2}, \tag{C.1}$$

where $\mathcal{F}_i^{\rm obs}$ are the observed fluxes from optical to gamma rays (*cf.* Figure 5.3), $\mathcal{F}_i^{\rm obs}$ are the predicted fluxes at the same frequency, and σ_i is the standard deviation of the measurement

[1] Available at http://github.com/DEAP/deap, as of the date of publication of this thesis.

distribution that yields each data point. For *Fermi* data, we used as standard deviation the size of each error bar; for the optical observations, we considered the spread of flux values observed during the period of the flare. Although technically the definition of χ^2 would only apply to a Gaussian distribution, which is not necessarily the case for all the measurements, we found that it provides a good estimate of the goodness of the predicted SED at a given wavelength for purposes of optimization. Regarding the X-ray constraint, since it is only an upper limit, we considered the following value:

$$\chi^2_{i=\mathrm{Xray}} = \frac{(\mathcal{F}^{\mathrm{model}}_{\mathrm{Xray}} - \mathcal{F}^{\mathrm{obs}}_{\mathrm{Xray}})^2}{\mathcal{F}^{\mathrm{obs}\ 2}_{\mathrm{Xray}}} \times \theta(\mathcal{F}^{\mathrm{model}}_{\mathrm{Xray}} > \mathcal{F}^{\mathrm{obs}}_{\mathrm{Xray}}), \tag{C.2}$$

which means that SEDs whose X-ray flux does not overshoot the observational constraint are not penalized ($\chi^2_{i=\mathrm{Xray}} = 0$).

Besides the requirement that the SED be in agreement with the observed fluxes, we also wish to predict a number of muon neutrino events in IceCube that is as close as possible to 13 (Aartsen et al., 2018c). We therefore add to our overall quality criterion a goodness-of-fit parameter of the neutrino emission, given by

$$\chi^2_\nu = \frac{(N^{\mathrm{model}}_\nu - N^{\mathrm{obs}}_\nu)^2}{N^{\mathrm{obs}}_\nu} = \frac{(N^{\mathrm{model}}_\nu - 13)^2}{13}, \tag{C.3}$$

where the variance σ^2 has been replaced by the observed number of neutrinos, drawing from Poissonian statistics.

We then obtain an overall fitness function of each individual of

$$f = \chi^2_{\mathrm{SED}} + \chi^2_\nu. \tag{C.4}$$

More precisely, the goodness of an individual is given by how small is its value of f. This way we have reduced the optimization of a complex system into a one-dimensional minimization problem.

Algorithm– The flow of the algorithm is represented schematically in the diagram of Figure C.1. The steps of the algorithm can be summarized as follows:

1) A population of N random individuals is created, each consisting of a list of seven parameters, distributed uniformly across a wide range of the parameters. We used a population size of $N = 10^4$. At a first stage, all the individuals are evaluated according to the above procedure. This implies running the lepto-hadronic code AM^3 ten thousand times, which could be done in parallel using the computing cluster at DESY Zeuthen[2].

[2]I would like to thank Leonel Morejon for useful discussions about the implementation of the algorithm.

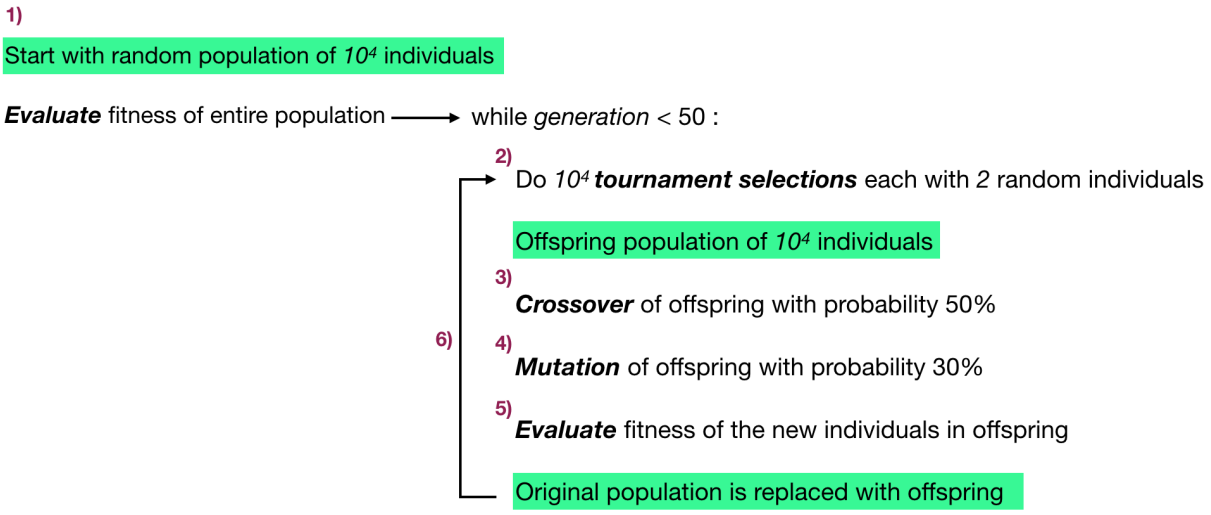

Figure C.1: Simplified schematics of the flow of the genetic algorithm used for the optimization of the physical parameters of TXS 0506+056 (Rodrigues et al., 2018b). The steps illustrated in the diagram are explained in detail in the text.

2) In the next step, the individuals compete to remain in the next generation based on their fitness functions. Until the end of the algorithm all generations will have the same number of individuals, so the selection process must always guarantee that N individuals proceed to the next generation; however, some may be selected more than once (increasing their presence in the next generation), while others may not be selected at all and be extinguished. Rather than making the entire population compete by selecting each individual with a probability based on its fitness, we use the indirect method of *tournament selection*. A set of $k = 2$ individuals is *randomly* chosen, and from that pair, based on the fitness values of the two individuals, one is selected to pass to the next generation. This process is repeated N times until Compared to this method, a direct competition within the entire population (equivalent to the case $k = N$) would lead to a fast convergence towards the best existing solutions; on the contrary, picking random pairs of individuals to compete helps maintain the *variability* of the population, therefore assuring that the parameter space is scanned thoroughly before the population converges (Miller et al., 1995).

3) Once the next generation (*i.e.* the *offspring*) has been selected, we now have again N individuals, all of which were present in the previous generation; the only difference is the frequency of each one in the overall population. We now need to inject variability into the offspring, in order to assure that the largest possible part of the parameter space is scanned. This is done first through the process of *crossover* of different parameters between pairs of individuals (akin to the biological process of genetic recombination, which provides the offspring

with a combination of genes from both progenitors). In this case we consider a two-point crossover, which means that a random segment of the list of parameters of one individual is exchanged with the same segment of the parameter list of its mate. Pairs of individuals are randomly chosen across the entire population, and the probability of crossover occurring between each pair is fixed to 50% throughout the simulation.

4) Another process that must be included in order to sustain the variability of the offspring is that of *random mutations*. For every individual, we change the values of its parameters with a fixed probability of 30% (for each one of the seven parameters the individual probability of mutating is $30/7 \approx 4\%$). When a parameter is mutated, its value is changed to a new value within the original range of that parameter. Compared to crossover, which can move an individual very far in the parameter space by switching a portion of its parameters with another, mutations provide mostly *local* variability, by stochastically generating new solutions that are mostly close to the original individual[3].

5) Once the crossover and mutations have been applied to the offspring, the individuals that have suffered changes will be re-evaluated by computing the emitted SED and neutrino spectrum and applying the fitness function Eq. (C.4). The offspring then replaces the original population.

6) Steps *2)* through *5)* are repeated until the quality of the best individuals eventually stagnates. In subsequent generations, the population gradually converges to the best sub-sets of the parameter space, and the number of re-evaluations will decrease, since the same "good" individuals will appear with higher frequency in the population. However, as noted previously, this convergence must be slow enough that the population does not fall into a local minimum without having properly probed the full parameter space. In this work, we have simulated 50 generations. The two examples shown in Figure 5.3 were obtained in two separate simulations, each with a different (fixed) value of accretion disk temperature $L_{\rm disk}$. The examples correspond to the best individual in the final generation of the respective simulation.

[3]The probability distribution function of the new parameter value is a polynomial (given by the function `mutPolynomialBounded()` of the DEAP toolbox, with the argument $\eta = 20$). The distribution is highest for values that are closer to the old value, similar to a Gaussian distribution centered around the old value, but is limited to the original range determined for that parameter, rather than extending to infinity.