The Radio Signal of Cosmic Ray Air Showers and its Synthesis through Templates

Das Radiosignal von Luftschauern kosmischer Strahlung und dessen Synthese mittels Templates

Zur Erlangung des akademischen Grades eines
Doktors der Naturwissenschaften
von der KIT-Fakultät für Physik
des Karlsruher Instituts für Technologie (KIT)

genehmigte
Dissertation

von David Alan Butler (M.Sc.)
aus Dachau, Deutschland

Tag der mündlichen Prüfung:	24.01.2020
Referent:	Prof. Dr. Ralph Engel
Korreferent:	Priv. Doz. Dr. Torsten Enßlin
Betreuer:	Dr. Tim Huege

CONTENTS

1

INTRODUCTION

Most observations of the universe are made through photons, from ancient stargazing to modern radio and gamma ray astronomy. But many other messengers are also ubiquitous and can provide important complementary information or insights into objects and phenomena which are completely opaque to electromagnetic radiation. The measurements of cosmic neutrinos [1] and gravitational waves [2] were recent historical milestones in part due to how exotic and difficult to detect they are. By comparison the measurement of electromagnetically interacting particles other than photons would appear to be a far simpler prospect at first. All such particles of astrophysical origin carrying an electromagnetic charge are termed cosmic rays (CR) due to historical confusion with high-energy photons. Additionally common source models for other high-energy messengers often include conditions capable of accelerating cosmic rays as well: x-rays and gamma rays can be produced via scattering processes in turbulent plasmas, secondary gamma rays and neutrinos have the greatest chance of escaping high-density target material struck by a primary beam of even greater quantum energy and the known sources of gravitational waves, neutron stars and black holes, are known to host incredibly powerful magnetic fields compressed and "frozen" during the collapse of their parent star. These magnetic fields then drive relativistic jets into the interstellar medium which may arrive on Earth as cosmic rays.

Accurate measurement of cosmic rays is not simple however, chiefly due to the strength and diversity of their interactions which establish them as primary acceleration targets to begin with. Further due to the low flux at the highest energies modern experiments instead rely on sparse arrays in order to save costs. An alternative approach is to use precise detectors and extract as much information as possible from each individual event. One such measurement method relies on secondary radio-wavelength radiation emitted as the primary cosmic ray interacts with the Earth's atmosphere. This radio component is tied to large-scale coherence effects which can largely be described using classical electrodynamics and also smooth out many of the fluctuations common in the preceding high-energy particle interactions which makes it lucrative for indirect calorimetric measurements. In addition to electromagnetic calorimetry dense radio arrays such as LOFAR [3][4] can measure the development of atmospheric cascades induced by cosmic rays with impressive precision [5] which can then be related to the properties of primary particle itself, in particular its mass for nuclei.

Herein lies a major challenge for the radio technique: while the immediate radio emission processes are believed to be quite well understood their relation to the primary cosmic ray is not. In order to accurately describe modern measurements we rely on computer simulations of the high-energy particle physics which can be quite demanding and prohibitively so for the most accurate techniques. For example inferring a single quantity of interest via its χ^2 distribution as in the cited LOFAR method requires hundreds of single-event samples. Fully analytical descriptions are fast and thus could become viable in the future but currently do not reproduce measured data with sufficient accuracy. In this work we set out to construct the core of a hybrid simulation model using accurate simulations as templates. These templates are then processed based on our analytical understanding of cosmic ray radio emission thus eliminating the computationally demanding phase from the end-user application.

Such an approach does face known theoretical limitations arising from quantum fluctuations in the interactions between cosmic rays and material in the detection volume. Due to the sheer number of secondary interactions many of these fluctuations amount to the same average under the correct parametrisations in a phenomenon known as *shower universality* rather than yield unique jet-like events as in the clean environment of the LHC. But this universality also has known limits and the universality of radio emission specifically has not been studied to the same level of detail as the particle interactions themselves. Our work thus includes both exploration of universality and its limits specifically for radio emission as well as the application of the acquired knowledge in the construction and refinement of our template-based synthesis model. If successful the completed model would enable the use of far more powerful inference methods to extract the defining properties of primary cosmic rays in measured events.

Both the fast forward simulation tool and the knowledge of radio universality are incredibly lucrative for next-generation dense radio arrays such as the SKA-low [6] as they offer the opportunity to study cosmic ray signals with unprecedented precision. By simple geometrical projection and information entropy arguments the SKA-low should allow direct interferometric imaging of the radio sources developing in the atmosphere or the application of dedicated statistical inference methods such as Information Field Theory [7] to directly extract as much information on the cosmic ray primary as possible. Both approaches depend on a fast and precise forward simulation model against which to query their predictions.

2

HIGH-ENERGY COSMIC RAYS AND AIR SHOWERS

2.1 COSMIC RAYS

In practice most cosmic rays are nuclei of elements known to be common in our universe while at lower energies we also observe free electrons. One obvious component of the incoming charged particle flux is the solar wind which gives us a chance to study potential emission mechanisms directly. At the surface of the sun charged particles are accelerated from the ambient plasma through interaction with strong magnetic fields which remains the basic source model for cosmic ray accelerators even up to the highest energies. This is represented by the Hillas criterion [8] as an estimate of any given source's maximum cosmic ray energy E_{max}

$$E_{\mathrm{max}} = q \cdot B \cdot \beta_{\mathrm{acc}} c \cdot R \tag{1}$$

where q is the charge of the cosmic ray, B the average magnetic field strength in the source, β_{acc} a characteristic acceleration strength, c the vacuum speed of light and R the physical size of the acceleration region. It should be noted that similar to modern particle accelerators the magnetic field B primarily serves to confine charged particles as they accumulate energy, the actual acceleration mechanism need not be driven by that same magnetic field. For example the model of shock acceleration [9] contains one very compressed electromagnetic accelerator in the shock front characterised by the shock velocity $v_s = \beta_{\mathrm{acc}} c$ and then the uncompressed ambient field B which confines the cosmic rays and drives them back and forth through the shock front, similar to the basic construction of a laser resonator.

Based on the Hillas criterion and supported by flux estimates as well as modern simulations many cosmic rays originate from known astrophysical phenomena such as stellar winds and supernovae. For these sources the acceleration limits are well known but still far below the highest observed cosmic ray energies which motivates the study of ultra-high energy cosmic rays (UHECRs) in particular. With individual particle energies far above the capabilities of man-made technology we need to turn to more exotic high-energy phenomena such as gamma ray bursts or jets from black holes, a selection of source candidates is shown in figure 1.

In contrast to the many forms of electromagnetically driven acceleration and confinement there are very few means of accelerating electrically neutral particles and explaining how they might arise naturally is even more challenging. Thus the preferred source models for neutral high-energy messengers

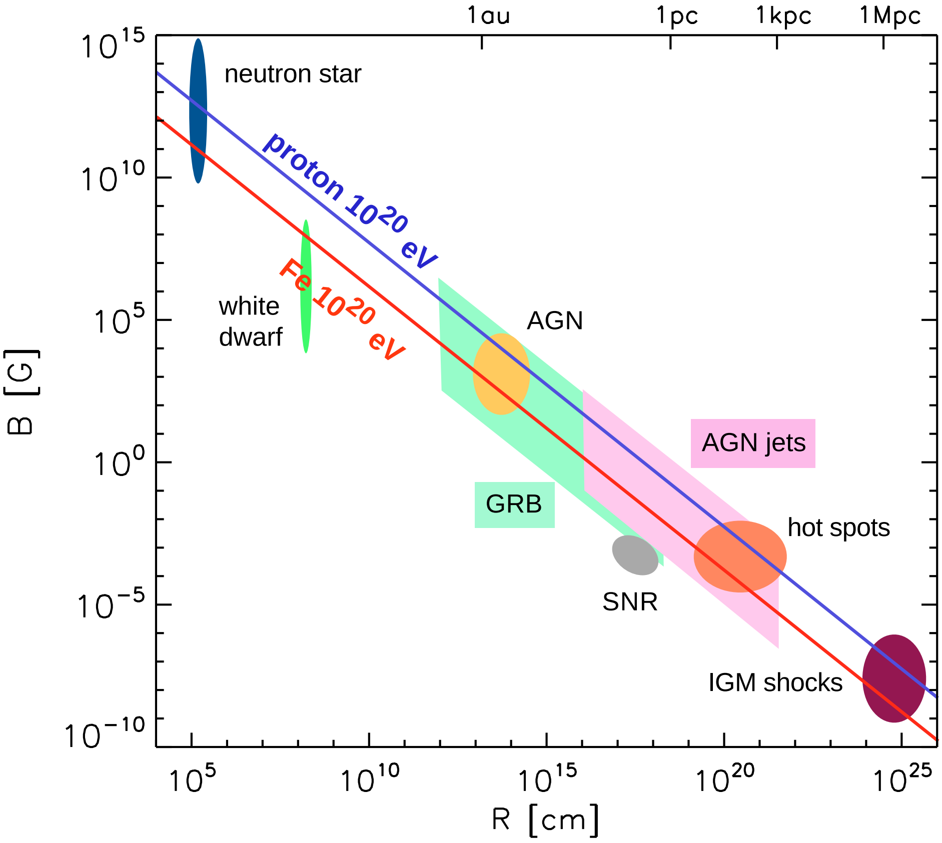

Figure 1: Hillas plot: double logarithmic graph of magnetic field strength over geometric size of the source. Objects above the red line can accelerate fully ionised iron nuclei above 10^{20} eV, those above the blue line can accelerate protons above 10^{20} eV. Source: [10]

such as gamma rays or neutrinos rely on a local population of charged particles attaining even higher energies in order to produce them as secondary particles in high-energy collisions [11]. As charged particles escaping from distant sources are deflected by large-scale magnetic fields they have very low chances of reaching us let alone pointing back to their source. On the other hand the detection of high-energy cosmic rays implies the existence of neutral secondary particles of which the most energetic should preferentially point towards the cosmic ray accelerators [12]. This is one of the key goals of modern multi-messenger and UHECR astronomy: to combine measurements of different particles in order to identify the sources of the most energetic particles in the universe.

Cosmic rays at the highest energies are quite rare however: with increasing cosmic ray energy their flux decreases significantly as shown in figure 2. At the GeV scale an observed flux over 1 $\mathrm{m}^{-2}\mathrm{s}^{-1}$ allows effective direct study with space-based particle detectors. In the PeV range with fluxes around 1 $\mathrm{m}^{-2}\mathrm{a}^{-1}$ such missions yield low statistics even from prohibitively expensive multi-year campaigns thus a different approach becomes necessary. Due to the rapidly decreasing flux we would still ideally like to perform full calorimetric measurements and extract as much additional information from every single event as possible, without massive space-based precision detector arrays this poses a significant challenge however.

2.2 COSMIC RAY AIR SHOWERS

On Earth incoming cosmic rays naturally strike our atmosphere resulting in a cascade of secondary particles. These particle cascades are called Extensive Air Showers (EAS) and can cover the entire height of the Earth's atmosphere while spreading out laterally over hundreds of metres by the time they reach solid ground making it possible to indirectly detect fluxes below 1 $\mathrm{km}^{-2}\mathrm{a}^{-1}$ with large ground-based detector arrays. The downside is that the exact outcomes of cascade interactions are determined by quantum field interactions meaning we can never deterministically reconstruct the full initial state. Further the initial stages can involve centre-of-mass energies above those accessible to modern accelerators adding significant systematic uncertainties in the interaction crosssections. Eventually the energy is spread out across a large number of secondary particles at which point ensemble descriptions can be reproduced in lab experiments at the cost of no longer being able to resolve individual particles without prohibitive detector costs. In principle the centre-of-mass energy is high enough to produce plenty of any known particle but in practice the branching ratios of the initial highest-energy interactions lead to distinct types of cascade depending on the primary particle:

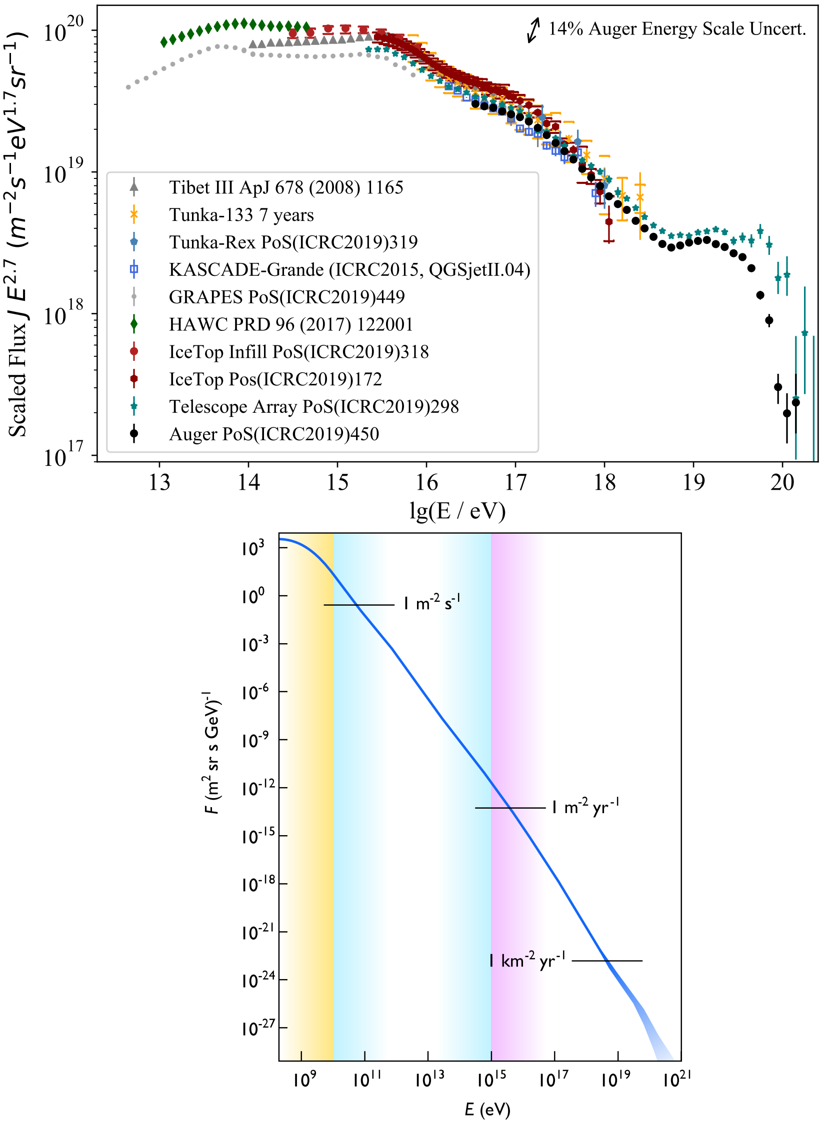

Figure 2: Top: spectrum of the high-energy cosmic ray flux as measured by several experiments both historical and ongoing. General acceleration physics predict a simple power law spectrum which has already been divided out using an exponent of -2.7. All remaining spectral features are thought to arise from a combination of propagation physics and overlap between different source populations. Source: [13]. Bottom: simplified schematic including the lower particle energies not observable in sparse indirect detector arrays. The lowest energies (yellow) are primarily dominated by the solar wind, intermediate energies (cyan) are attributed to Galactic sources and cosmic rays at the highest energies (magenta) most likely are accelerated outside our Galaxy. Source: [14]

- **Electromagnetic Cascades:** electrons, positrons and gamma rays greatly favour electromagnetic interactions producing entire cascades dominated by these three particle species [15]. Energy is propagated downwards in a fairly regular fashion thus simple mathematical models suffice to capture the macroscopic behaviour of the air shower. Due to the limited range in air only a declining population of low-energy secondary particles commonly reaches sparse ground-based detectors.
- **Hadronic Cascades:** Hadrons such as the nuclei representing virtually all known cosmic rays at the highest energies [16] additionally interact via the strong and weak nuclear forces. Similar to detectors in collider experiments this leads to jets and sub-showers which cause local macroscopic overabundances not treated by basic cascade theory. Ultimately the decay of secondary pions as well as ionisation losses and bremsstrahlung yield a significant electromagnetic component as well thus allowing the use of identical detection techniques. Additionally however the decay of charged pions produces a significant number of muons which also reach solid ground and thus offer an additional measurement channel [17].
- **Exotic Events:** While not relevant for this work it should be mentioned various rare processes do not fully fit the above categories. Neutrinos as the most widely studied outlier can penetrate deeply into the atmosphere and even solid ground which, despite vanishingly small cross-sections, could induce a cascade originating close to ground level or in the shadow of mountains where cosmic rays or gamma rays would be unable to penetrate [18].

Neutrinos themselves are not included in the definition of cosmic rays however the charge of the primary particle cannot be measured directly from air-shower cascades and even if it could the most common neutrino interactions produce a charged lepton carrying most of the initial energy. For this reason modern neutrino detectors often work in close collaboration with cosmic ray measurement efforts. In this work we focus on hadronic cosmic rays and thus the more complex cascades.

2.3 ATMOSPHERE AND CASCADE QUANTITIES

Mathematical cascade theories operate in abstract dimensionless parameters. In order to represent physical reality we need to map this abstract space to measurable quantities, most importantly the evolution of the cascade itself. Initially the primary particle effectively travels in a straight line with constant velocity $\vec{v}$. This line of propagation is called the *shower axis* and the distance reference is usually chosen to be the location of the first interaction in the atmosphere. In detection and analysis one might also use the *shower core* defined as the intersection of the shower axis with the ground level of the detector. Due to the massive kinetic energy of the primary particle its velocity is effectively equal to the speed of light in a vacuum which yields a simple equivalence between traversed distance and elapsed time. This gives rise to the definition of the *shower front* as a surface expanding with velocity c from the point (and time) of first interaction, the causal horizon preceding all physical particles in the cascade. As most high-energy secondary particles remain highly relativistic physical timings within the cascade are often given relative to the imaginary shower front rather than the global time reference.

Cascade evolution is determined by particle interactions however. The most significant scale in this regard is the ambient density of the air serving as the primary target material. By abstracting our description to exclude the physical atmosphere we can apply simpler mathematical models and then re-add a detailed atmosphere later to map back to physical reality. In line with our picture of a steadily evolving cascade we use the atmospheric (column) depth X as our fundamental parameter, integrating over the local density

$$X(h) = -\int_{\infty}^{h} \rho(h)\,\mathrm{d}h \tag{2}$$

which for a standard exponential density profile $\rho(h)$ is fully invertible. In analyses of real measurements one often uses the slant depth X_{slant} along the shower axis $\vec{l}$

$$X_{\text{slant}}(l) = -\int_{\infty}^{l} \rho(\vec{l})\,\mathrm{d}\vec{l}. \tag{3}$$

Specific to the vertical showers we perform most tests with slant depth and vertical column depth are identical and the curvature of the atmosphere plays little role. Nevertheless the exponential scaling of the atmospheric density still factors into the relation between atmospheric depth and geometrical height, distance or propagation time.

Our simplest summary of an air shower cascade is the longitudinal profile, the one-dimensional evolution of the total number of particles over atmospheric depth. Traditionally one counts all charged particles however in theoretical modeling we will also use the profiles of specific subgroups of particles such as hadrons or only electrons and positrons. As shown in figure 3 the longitudinal profile matches a bell curve, upon closer inspection it is not a simple Gaussian however. The additional features are described by the so-called Gaisser-Hillas curve [20] which has since been refined through additional parameters to the form in use today

$$N(X) = N_{\text{max}} \cdot \left(\frac{X - X_1}{X_{\text{max}} - X_1}\right)^{\frac{X_{\text{max}} - X_1}{a + bX + cX^2}} \cdot \exp\left(\frac{X_{\text{max}} - X}{a + bX + cX^2}\right) \tag{4}$$

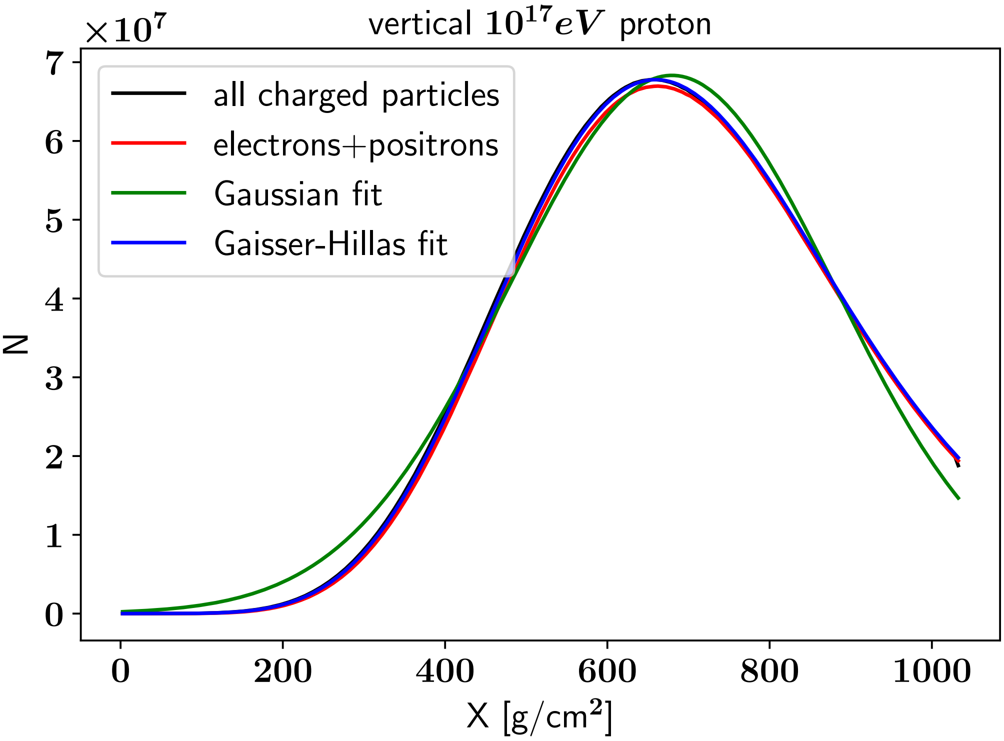

Figure 3: Longitudinal profiles of a simulated [19] vertical extensive air shower initiated by a proton with a primary energy of 10^{17} eV. The Gaussian (green) and Gaisser-Hillas (blue) fits were calculated with respect to all charged particles. Electrons and positrons (red) dominate the raw number of particles throughout the entire evolution of the air shower.

where $X_{\max}$ describes the position of the curve's maximum and $N_{\max}$ the number of particles at said maximum. X_1 represents the location of the first interaction between the primary cosmic ray and a particle in the atmosphere. The quadratic polynomial of a, b and c originated as a scalar interaction length X_0 akin to the mean free path length often used in the microphysics of ensembles. All of these parameters are fitted to the data of each individual cascade although empirically the location of the shower maximum, $X_{\max}$, has proven the most useful by far in reconstruction and analysis to the point of often being abbreviated as "the shower maximum". Efforts to reproduce measured particle distributions from the longitudinal profile and simple statistical models were unable to maintain pace with more accurate detectors leading to the development of complex computer models.

2.4 MICROPHYSICS SIMULATION - CORSIKA MONTE CARLO CODE

Our most successful numerical models completely forego the macroscopic picture in favour of individually evaluating the interactions between microscopic particles [19][21]. With interaction energies far into the realm of quantum field theories these simulation frameworks resemble the design of Monte Carlo sampling techniques in statistics albeit with all inputs beyond the numerical random number generator itself informed by physics to the best of our ability. This approach far outperforms all existing competitors in accurately reproducing measured data [22][23] cementing it as the *de facto* standard in analyses and our best representation of the truth. Nevertheless the empirically determined input paramters such as interaction crosssections and atmosphere models themselves constantly undergo refinements and corrections leading to degree of systematic uncertainty beyond the random fluctuations and limited sampling inherent to the methodology of the simulation framework itself. This compounds the primary drawback of basing entire analyses on numerical Monte Carlo simulations: each individual shower requires substantial computation time and with no hierarchy of scales or analytical understanding of the dynamics all simulations must be recomputed following any revision of the code in order to update previous analyses.

Despite these challenges the cosmic ray community constantly compares different models and numerical implementations with generally promising results. In particular the dynamics of the electromagnetic component relevant to this work tend to be quite consistent [24][22]. Thus studying the impact of our choice of input models will be left to later verification and these systematic uncertainties will not factor into our evaluation.

The particular program used in the context of this work is called **CO**smic **R**ay **SI**mulations for **KA**scade [19] though it is not specific to the Kascade [25] experiment any longer. The specific program versions used in our simulations range from 7.5000 to 7.6400 with some retroactively applied bug fixes to the earlier versions. For our high-energy interaction model we chose QGSJETII-04 [26] as bundled with the CORSIKA release and for low-energy interactions we used the closed-source FLUKA library [27]. All other parameters are effectively handled by CORSIKA at the runtime of each individual air shower simulation including the choice of atmosphere and may vary, although most simulations were performed with different parametrisations of the US Standard atmosphere [28].

Simple cascade models still remain in use when only broad estimates or averages are needed and when integrated into the modern simulation tools they can also provide preliminary estimates or partial screening of the results. Of particular note to this work is the integration of CONEX [29] into the CORSIKA code framework allowing the rapid (albeit somewhat inaccurate) prediction of macroscopic properties. Rather than act as a standalone mathematical model CONEX can receive the secondary particles produced by initial high-energy interactions in CORSIKA as its input and continue with analytical cascade equations. This constitutes our best means of predicting the shower maximum of a CORSIKA simulation without investing the computation time of a full simulation by noting down the numerical random seeds of the simulation input. When a shower maximum within the desired range is found via the much faster sampling in CONEX a full CORSIKA simulation can be started with identical input parameters and usually matches the predicted depth of the shower maximum within 5-10g/cm^2 for our configuration.

3

RADIO EMISSION FROM COSMIC RAY AIR SHOWERS

As the discourse so far implies the central component to cosmic ray air shower studies always has been the high-energy particle physics starting with the discovery of "Höhenstrahlung" - ionising radiation increasing in intensity with increasing height above solid ground - and continuing to this day with most large-scale detectors such as Kascade [25] and the Pierre Auger Observatory [30] being designed around a core array of particle detectors. As the yield of secondary particles at ground level diminishes greatly with increasing distance to the air shower and remains prone to intrinsic fluctuations both in production and in direct detector types the community constantly studies potential enhancements. One very successful approach involves combining direct measurements of air shower particles with secondary signals generated by the cascade. These can appear in vastly different parts of the electromagnetic spectrum including Cherenkov radiation [31] due to the high-energy core of the air shower traveling faster than the speed of light in air, fluorescence of air molecules excited by the high-energy radiation [32] and the component this work focuses on: coherent radio-frequency signals generated by the large-scale dynamics of charged particles within the air shower.

Indeed radio signals offer several key advantages. They propagate with next to no attenuation through air and even ice [33] diminishing in amplitude mostly according to flux conservation laws. As the coherent radio emission is completely dominated by the contributions of electrons and positrons it can be used to sample the electromagnetic component of the air shower. When comparing integral quantities such as the total amount of energy the radio emission deposits into the ground it turns out the complex cascade dynamics have fairly little influence as they primarily redistribute the emitted energy to different positions at ground level and the very large number of secondary electrons and positrons smooth out local fluctuations in the high-energy hadronic interactions. In summary the energy in the entire radio footprint at ground level can be reliably related to the total energy in the electromagnetic component of the cascade to obtain a calorimetric measurement.

While optical- or near-optical-wavelength fluorescence detectors could capitalise on isotropic emission they are completely blinded by the sun or even strong moonlight as well as cloud cover leading to a very low viable duty cycle [34]. Radio antennas meanwhile can operate under almost all conditions only being compromised by thunderstorms. In addition radio-frequency hardware generally is cheap to produce and deploy when compared to particle detectors or optical telescopes and many challenges are well understood from its varied applications in telecommunication. Therein also lies one of the greatest pitfalls of the radio technique: many electronic technologies and *especially* modern wireless communications create intolerable levels of background noise and even extremely short pulsed signals potentially indistinguishable from cosmic rays.

While modern data analysis techniques are beginning to overcome the challenges of interfering signals at the time of writing radio measurements have only truly proven themselves as an additional component to particle detector arrays whose event triggers do not respond to purely electromagnetic signals. Thus the main goal of radio measurements is to provide as much supplementary information as accurately as possible which requires a sound understanding of the signal beyond numerical simulations.

3.1 CLASSICAL MACROSCOPIC THEORY

Indeed theoretical models were proposed long before radio signals from cosmic ray air showers were actually measured. First predicted by G. Askaryan in 1962 [35] the ionisation of the atmosphere was expected to accumulate a net negative charge within the air shower cascade capable of emitting electromagnetic radiation. Despite informing initial discovery experiments the predicted Charge Excess (CE) or Askaryan emission actually turned out to be subdominant. The majority of the observed radiation originates from interactions with the Earth's magnetic field $\vec{B}$ [36] and is thus called Geomagnetic radiation (Geo). As both mechanisms can neither be observed individually in cosmic ray air showers nor easily extracted from Monte Carlo simulations their characteristic polarisation patterns as well as the Geomagnetic component's direct scaling with the ambient magnetic field strength are considered their defining properties.

For this reason radio analyses like to define a local coordinate system by the shower axis $\vec{v}$, the alignment of the primary Lorentz force $\vec{v} \times \vec{B}$ and the third orthogonal axis $\vec{v} \times \vec{v} \times \vec{B}$. Meanwhile our primary lab frame is chosen to be consistent with CORSIKA [19] meaning our $\hat{x}$ axis points north in rough alignment with the geomagnetic field, the $\hat{y}$ axis points west and the $\hat{z}$ axis points up while all air showers we consider propagate *downwards*. An overview of both coordinate systems is shown in figure 4. Signs may be omitted or adjusted for visibility meaning axes of matching direction but different orientation may be used interchangeably but both coordinate systems formally are right-handed. Care should be taken when comparing to literature as some descriptions may rotate the lab frame in increments of 90° or use distances relative to a particular detector at ground level.

Regardless of the exact emission process any mechanism driven by the cascade dynamics innately aligns with the co-moving frame of reference only

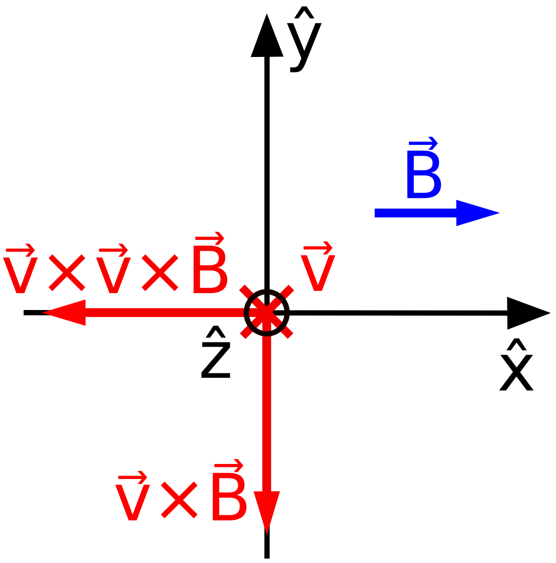

Figure 4: Coordinate systems used in this work for the example of a vertical shower: the lab/simulation frame $(\hat{x}, \hat{y}, \hat{z})$ is shown in black, the shower frame $\left(\vec{v}, \vec{v}\times\vec{B}, \vec{v}\times\vec{v}\times\vec{B}\right)$ in red and the ground projection of the geomagnetic field in blue.

slightly behind the shower front. Projected back into the observer frame even isotropic emission is concentrated into a cone facing forwards from the direction of propagation in a process known as Relativistic beaming or Doppler beaming. The geometric component of this transformation involves the velocity of the radiation which depends on the refractive index as c/n. Due mathematical similarity to the kinematics of the famous Cherenkov radiation we refer to this mechanism as *Cherenkov beaming* with its *Cherenkov Angle*

$$\theta_C = \arccos\left(\frac{1}{n}\right) \tag{5}$$

identifying the direction of the highest amplitude. It must be noted however that radio emission from cosmic ray air showers does **not** originate from the same process as the well-known blue glow in nuclear reactors. Further we observe a second-order beaming effect based on coherence along the shower evolution. This occurs because the refractive index is not constant. The refractivity $n-1$ scales with the local ambient density and thus follows an exponential profile in the vertical direction. Per equation 5 the (local) Cherenkov angle *increases* as the shower penetrates deeper and approaches the detector thus focusing more radio waves onto the same geometric distance in the observation plane at ground level despite the beaming *angle* constantly changing. We call this pronounced feature the *Cherenkov ring*.

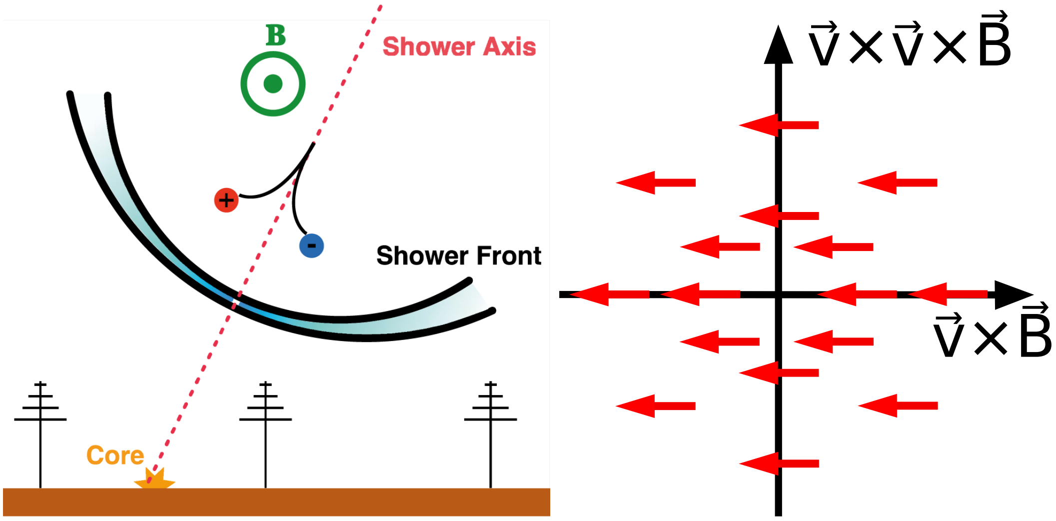

Figure 5: Emission mechanism (left) and polarisation pattern (right) of the Geomagnetic signal component. By definition the plane showing the polarisation pattern is always orthogonal to the shower axis and thus not necessarily parallel to even ground. Source for left graphic: [37]

3.1.1 *Geomagnetic Mechanism*

Being nearly homogeneous at these scales the geomagnetic field drives coherent charge separation as positively and negatively charged particles are deflected in opposite directions. While all charged particles are subject to this geomagnetic force the effective separation is limited by the mass of the particles. Combined with their sheer number this leads to electrons and positrons completely dominating the mechanism in practice. Rather than accelerate constantly frequent deflections in collisions with the ambient medium lead to a relatively constant macroscopic drift velocity. The resulting electric current driven by the Lorentz force can be pictured similarly to a Hall current in solid conductors only on vastly different time and density scales. Not radiative in itself the emission of electromagnetic signals including a coherent radio-frequency component originates from the time variation of this current as the evolving cascade and thickening atmosphere constantly modulate the injection and absorption rates respectively.

From this macroscopic picture we can qualitatively predict the resulting radio signal. Its polarisation is expected to be homogeneous and orthogonal to both the geomagnetic field and the driving particles' direction of travel which on average still coincides with the shower axis. In local coordinates we expect all Geomagnetic emission to be polarised along the $\vec{v} \times \vec{B}$ axis for all potential antenna positions as illustrated in figure 5.

Fueled by the high-energy particles close to the shower axis with only slower scattering and drift perpendicular to it we expect the amplitude of the Geomagnetic signal to be radially symmetric around the shower axis. As many electrons and positrons are produced and consumed on or close to the axis itself the amplitude will be non-zero at the centre of symmetry although for the relevant frequencies we still observe its maximum in accordance with the predicted Cherenkov beaming. Sample time series for different distances along the same lateral axis are shown in figure 6.

3.1.2 *Charge Excess Emission*

Askaryan's original prediction was based on the lighter electrons being swept away within the propagating cascade while the heavier cations lagged behind. This effectively creates a time-varying cloud of negative charge in the shower front as ionisation by the cascade and the atmosphere again constantly change the magnitude of the electron excess.

The resulting polarisation pattern points radially inwards towards the shower axis, vaguely reminiscent of a time-varying point charge. This leads to coherent partial interference between the Charge Excess emission and the Geomagnetic component along the $\vec{v} \times \vec{B}$ axis however the $\vec{v} \times \vec{v} \times \vec{B}$ axis will always receive only a Charge Excess contribution. Most importantly placing an antenna on said axis separates the two components into orthogonal polarisations allowing direct simultaneous and independent access. This is particularly relevant in theory and modeling because our macroscopic prediction does not *a priori* imply time coincidence, a slight offset would effectively result in measuring circular polarisation contributions which present a far greater challenge. Indeed a small component of circular polarisation has been measured with a high-precision instrument [39] though it is usually neglected. However even in higher-order macroscopic predictions some effects may yield a small Charge Excess contribution in the $\vec{v} \times \vec{B}$ direction, for example the excess of electrons also experiences geomagnetic deflection and thus shifts off the shower axis.

Like the Geomagnetic signal we expect the amplitude of the Charge Excess emission to be radially symmetric. Unlike the Geomagnetic component it drops to zero towards the shower axis thus always producing a peak even when Cherenkov beaming alone would not be sufficient. In air and for typical geometries the Charge Excess emission only contributes about 5-20% of the full signal as shown in figure 8 and verified by measurements [40] allowing rudimentary analyses to even neglect it entirely however unlike the Geomagnetic component it is always present. Cosmic rays arriving parallel to the Earth's magnetic field at the detector location produce no Geomagnetic emission and could therefore be used as a direct test of theoretical predictions.

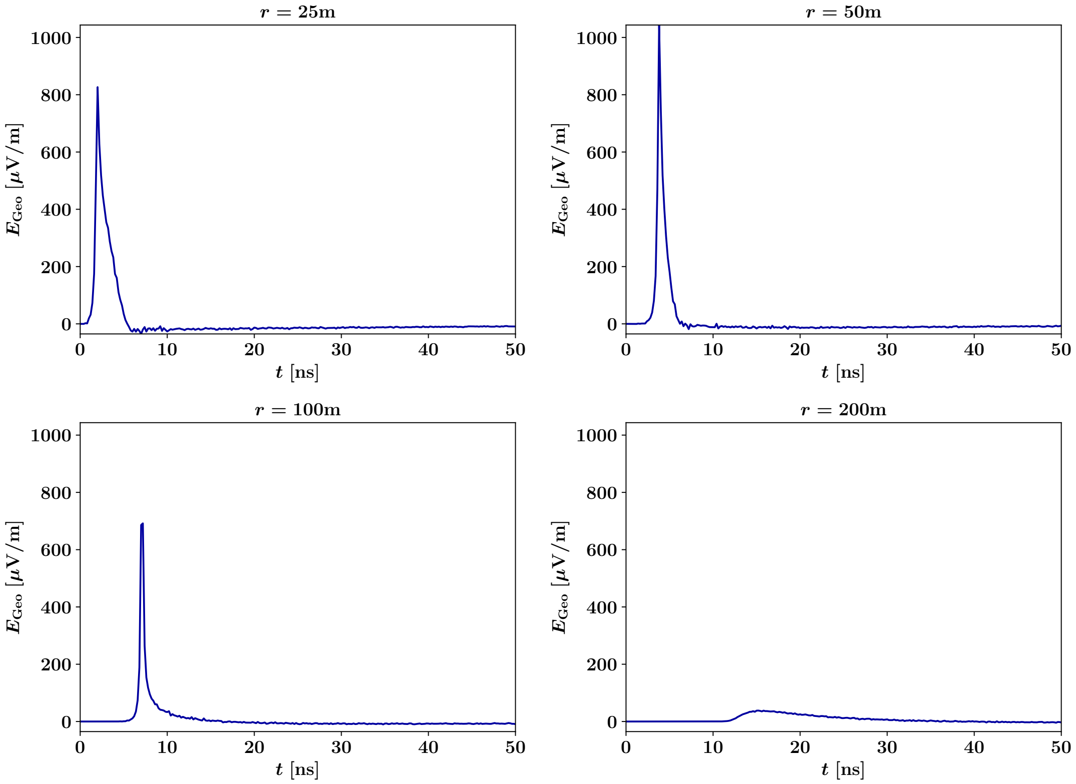

Figure 6: Sample time series of Geomagnetic emission for four lateral distances as simulated by CoREAS [38]. The antennas are positioned on the positive CORSIKA $\hat{x}$-axis corresponding to the *negative* $\vec{v} \times \vec{v} \times \vec{B}$ direction in the shower plane. Despite the extremely short duration of the pulses one can see slight differences in their shape and arrival times by eye. For large lateral distances the differences become quite drastic but the significantly lower amplitudes greatly hamper the detection of these pulses in real measurements. While none of the sample positions was chosen to match the theoretically predicted position of the Cherenkov ring the maximum amplitude is clearly reached between 25 m and 100 m rather than in the centre of the footprint.

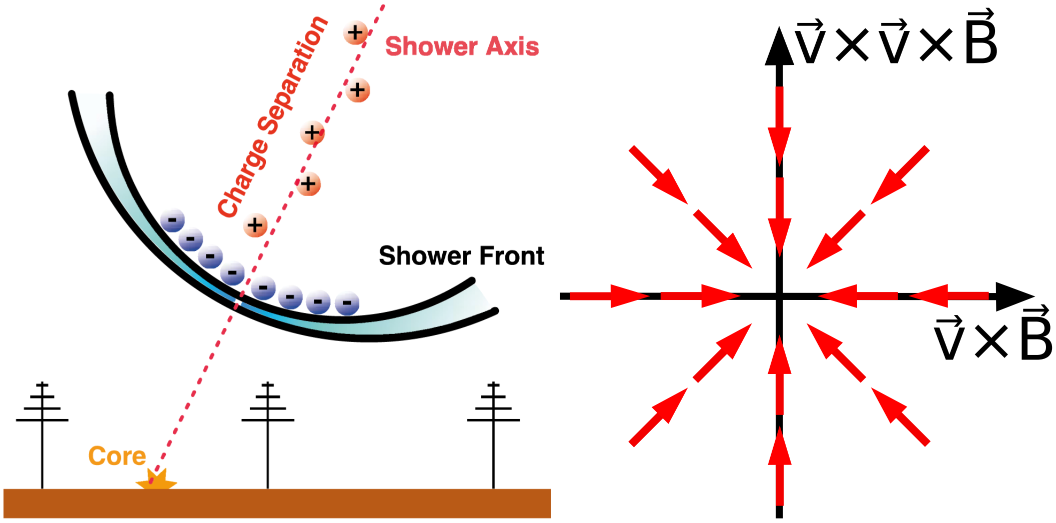

Figure 7: Emission mechanism (left) and polarisation pattern (right) of the Charge Excess signal component. Source for left graphic: [37]

3.1.3 *Digital Calculation in MGMR3D*

Representing this classical theory of radio emission from cosmic ray air showers we will be using the **M**acroscopic **G**eo-**M**agnetic **R**adiation (MGMR) model [36] or more precisely its expanded successor called MGMR3D [41]. In this framework the air shower is translated from a quantum particle cascade into a classical four-current which is then analytically parametrised via smooth functions. The amplitude of the four-current is explicitly split into the electrostatic charge expected to drive Charge Excess emission and the spatial drift current representing the Geomagnetic mechanism. These components as well as some spatial structure terms assumed to be common to both are then fitted to match average or common values observed in CORSIKA simulations. A schematic overview of the original MGMR model describing only Geomagnetic emission is shown in figure 9. While MGMR3D treats the charge excess independently in some respects it still maintains the picture of a common "shower pancake" speeding along the shower axis at almost the speed of light. The current $\vec{j}$ and charge distributions not only have a lateral structure but are also parametrised with respect to the geometric lag h behind the shower front in order to capture more of the true complexity of the air shower cascade.

It is important to note that despite extracting parameters to match CORSIKA during the design phase the MGMR model does not explicitly copy the radio signals produced by the Monte Carlo code. Instead the electromagnetic field at the location of the antenna is computed from the parametrised four-current following classical electrodynamics. As such it could, in theory at least, be calibrated using aerial measurements of real particle distributions

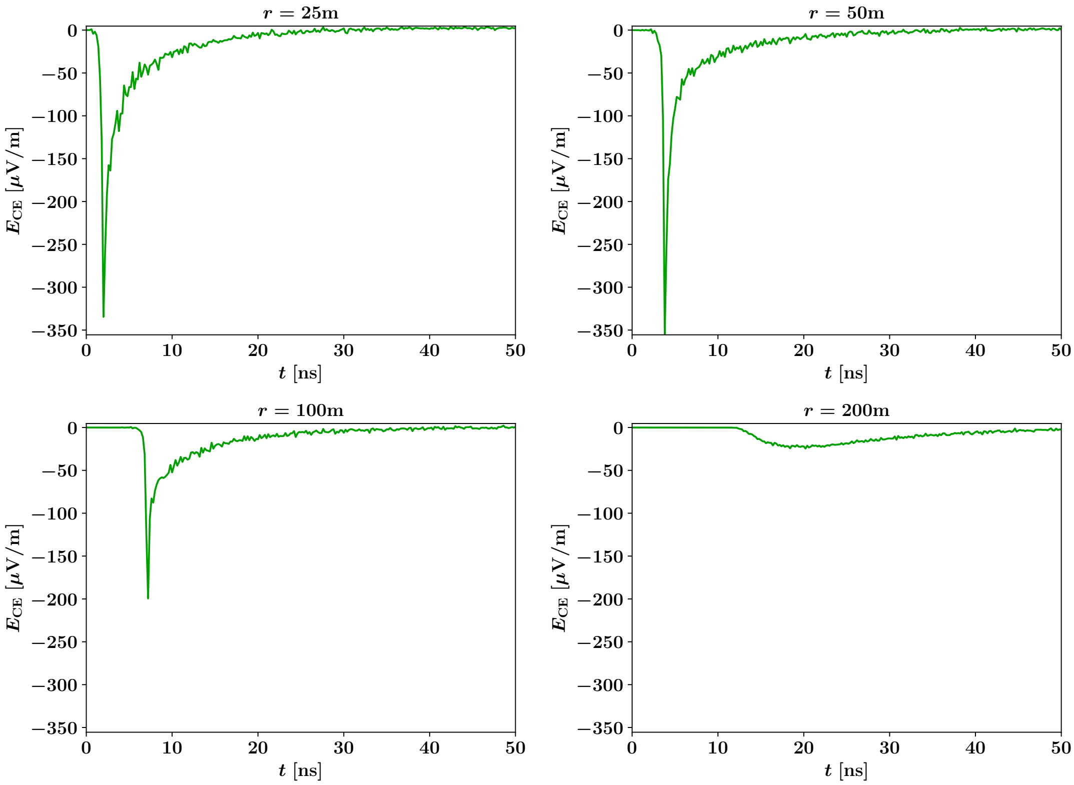

Figure 8: Sample time series of Charge Excess emission from the same shower and antenna configuration as used in figure 6. The time axis is identical but the amplitudes are significantly lower. Similarities in the lateral development between both polarisations hint at a common origin in propagation geometry while the differences in pulse shape are more likely to differentiate the actual emission mechanisms themselves.

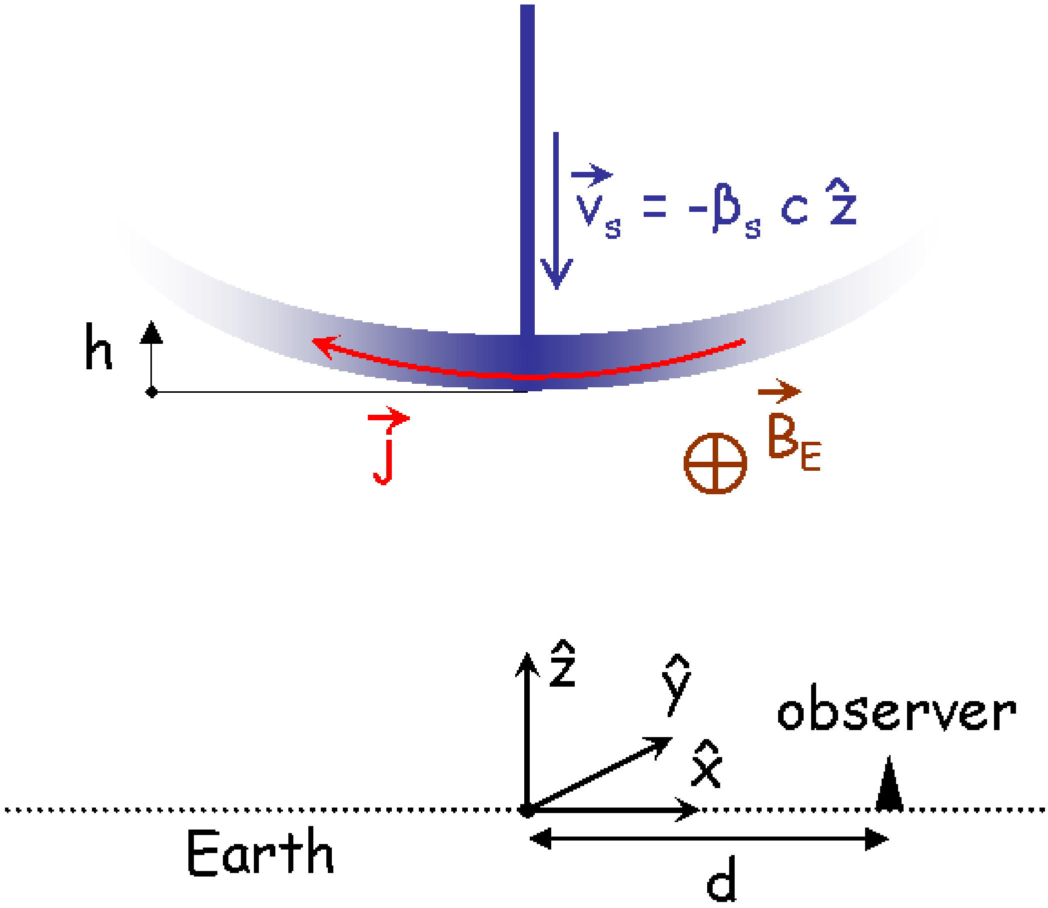

Figure 9: Schematic of the analytical MGMR source model. The shower is abstracted to a diffuse "pancake" (blue) representing relativistic secondary electrons and positrons. As this charged cloud propagates along the shower axis $\vec{v}_s$ (also blue) the Earth's magnetic field $\vec{B}_E$ (brown) induces a steady drift current $\vec{j}$ (red) of continuously varying magnitude. The lateral distribution and lag h of the electrons and positrons behind the shower front further modulate the current field adding complexity to the emitted radio waves. The radio-frequency signal is sampled at discrete observer positions on a line in the lab frame and then interpolated or extrapolated to match any specified observer position. Source: [36]

with no accompanying radio signals. One significant exception in the present version is the absolute scale of the output which is only given in arbitrary numerical units. Comparing to signal strengths derived from Monte Carlo simulations which do match the scale of measurements suggests a numerical MGMR3D electric field unit to roughly correspond to 3 nV/m. As there is no "official" scale calibration at the time of writing we will detail our own purpose-optimised approach for our application of the MGMR3D code in section 9.2.

3.2 MONTE CARLO SIMULATION OF RADIO EMISSION

As radio measurements became more accurate even the best macroscopic analytical models failed to predict all features observable in the data. The pragmatic alternative lies in applying classical electrodynamics to particles simulated in Monte Carlo programs which was already being developed concurrently with analytical models [42]. Early versions such as REAS relied on collecting aggregate ensemble data as an intermediate step which also proved inadequate, the true improvement came from treating microparticles individually thereby calculating the radio signal directly from independently verified particle dynamics. In addition to the inherent cost in computing time this approach poses the challenge of covering not only the phase space of observable particles but also that of possible antenna locations. Continuous fields are inherently challenging to represent in numerics thus the first clear simplification lies in matching detector architecture by only sampling the electromagnetic field at discrete positions in space. This concession remains inherent to all models for the forseeable future but in Monte Carlo-based codes it incurs the additional cost of needing to recalculate the entire particle cascade every time additional samples are required.

Beyond this fundamental limitation the exact specifications vary between different code frameworks some of which are inherited from the baseline particle cascade generator. As we use CORSIKA in this work our radio simulations all originate from the integrated **Co**rsika-based **R**adio **E**mission from **A**ir **S**howers (CoREAS, [38]) add-on. Another alternative which continues to show good agreement with CoREAS results [43] is available in ZHAireS (ZHS+AIRES, [44]) building upon the AIRES air shower generator. We do not explicitly compare our present work with ZHAireS but acknowledge its contribution as an independent verification tool.

3.2.1 *Endpoint Formalism*

One crucial component of CoREAS in which it differs from ZHAireS is the manner in which contributions from individual particles are calculated. Both start from the Liénard-Wiechert potentials valid for relativistic point charges but where ZHAireS sums up the individual potentials CoREAS first calculates the electric fields produced by each individual charge and approximates the result assuming far-field geometry. Next CoREAS implements the so-called *endpoint formulation* [45] in which the fundamental components represent a particle instantaneously accelerating from rest to an arbitrary velocity or vice versa where ZHAireS builds upon infinitesimal tracks of motion at constant velocity. Since CORSIKA itself communicates tracks as endpoints this allows CoREAS to remain consistent even when we apply unphysical selections or repurpose track data later in this work.

Ultimately both approaches should yield mathematically identical results as long as their assumptions remain valid and neither fully includes near-field contributions.

3.2.2 *CoREAS Code*

Beyond its direct application in interpreting measurements CoREAS offers substantial opportunities for extracting additional information. Sampling points or virtual “antennas” can be placed anywhere, even in the sky or underground, and receive the signal they should in open air. Further since particles are processed individually they can be pre-selected or binned based on certain characteristics granting access to knowledge not obtainable in lab experiments. These include the identity of the particle, its Lorentz factor and, crucial to this work, the atmospheric depth (or position) of the particle. This slicing of the particle ensemble allows us to map out which stages of the evolving cascade contribute to features in the radio emission, and by modifying the CoREAS code we can compute equivalents to macroscopic radiation sources such as charge densities or currents reminiscent of the older REAS approach.

Care must be taken not to fall prey to numerical limitations and artifacts however. The endpoints received by CoREAS usually represent numerical aggregates rather than real particles which can artificially enhance the coherence of their emitted radio signal or create excessively discretised results if too few contributions enter any particular selection. Sharp boundaries may also act as unphysical radio sources in themselves under the endpoint formalism if two destructively interfering contributions are assigned to different sides.

3.3 PULSE FILTERING

With our settings the raw output data of CoREAS contains a 400 ns trace with a sampling rate of 0.2 ns. This corresponds to a maximum frequency of 2.5 GHz in the frequency spectrum. The frequency spectra of CoREAS time series show increasing destabilisation or incoherence past a frequency range of about 300-500 MHz along with a steep drop in amplitude. While these features may originate from real physics they could also be related to numerical thinning or finite resolution effects. As no existing radio air shower arrays can accurately measure at such high frequencies we opted to remove them via a bandpass filter. The current benchmark goal is to work with the frequency range from 0-500 MHz which should be broader than all ranges sampled by existing and currently planned detectors. For practical tuning to existing experiments narrower bands can be considered to more closely match their sensitivity, for example the commonly used 30-80 MHz band is almost guaranteed to eliminate all numerical noise and most incoherent radio contributions.

For most lateral distances the unfiltered time series show an almost instantaneous rise of the signal followed by a fast decay, often resulting in total pulse widths below 10ns as shown in figures 6 and 8. This very narrow time interval causes major distractions when analysing pulse shapes as it overemphasises single-bin fluctuations and rounding errors. By comparison omitting higher frequencies not only smooths the time series but broadens its peaks as seen in figure 10.

A rectangular bandpass filter introduces artifacts through the discontinuities at the rectangle borders such as acausal "ringing" preceding the true arrival time of the pulse. As such this filter is not intended for practical application. Instead the intent is to limit our view and numerical evaluation of time series to frequency ranges relevant to experiments and to potentially simulate the loss of information.

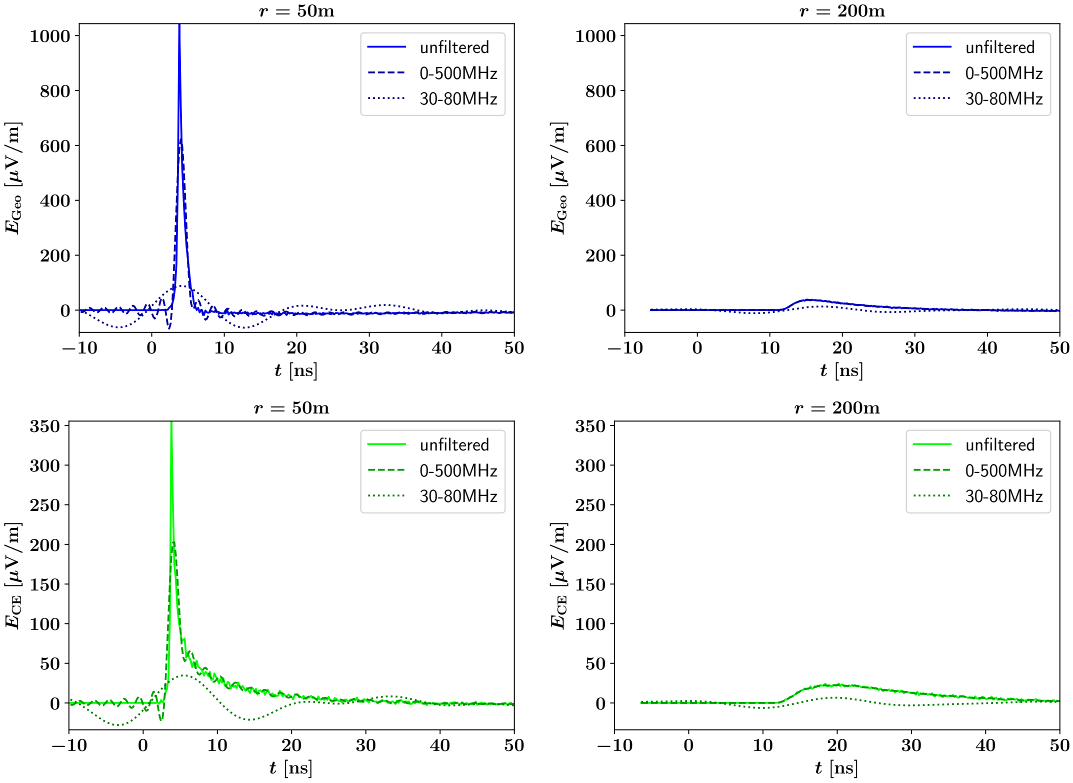

Figure 10: Sample time series of Geomagnetic (blue) and Charge Excess (green) emission for two lateral distances simulated by CoREAS [38] as already shown in figures 6 and 8. Additionally each time series is filtered to two frequency bands illustrating the smoothing and broadening effects of the filter as well as the "ringing" adding multiple zero crossings. The acausality of this ringing is not relevant to our work and subsequent application of a realistic antenna response will remove it.

4

UNIVERSALITY OF AIR SHOWERS

As already alluded to in section 2 the basic properties of electromagnetic air showers can be described by mathematically simple cascade models. One important discovery is the fact that even for hadronic cascades many aggregate quantities such as particle momentum distributions closely match simple empirical parametrisations. These models only depend on very few input parameters despite describing most instances of the highly probabilistic shower physics hence this observation is called *Shower Universality*. Reliance on universal behaviour forms the basis of many reconstruction techniques as it almost eliminates the possibility of complex correlations between multiple unknown quantities of interest. On the flip side in order to reconstruct physical properties of interest one must find observables which are *not* universal with respect to that property (but ideally to all others).

4.1 CASCADE UNIVERSALITY

Following Lafebre [47] the electromagnetic component, specifically the electrons and positrons which also happen to generate almost the entirety of the radio emission, is highly universal in most aspects of its phase space. They explain a higher degree of universality can be achieved by quantifying the longitudinal evolution of the cascade by its *Relative Evolution Stage* t_{RES}

$$t_{\mathrm{RES}} = \frac{X - X_{\mathrm{max}}}{X_0} \tag{6}$$

with $X_0 = 36.7\mathrm{g/cm}^2$ a constant approximation to the radiation length of electrons in air. Comparable performance is achieved with the more abstract *Shower Age* s

$$s = \frac{3X}{X + 2X_{\mathrm{max}}} \tag{7}$$

which approximates the parameter used in mathematical cascade theory. The important conclusion when comparing these quantities to the lab frame is the crucial role of the shower maximum with X_{max} providing the fluctuation-dependent adjustment.

The accuracy of the Gaisser-Hillas curve in fitting the longitudinal profile with only 3-6 free parameters as per section 2.3 can be understood as an early observation of universality. Dedicated studies go much further however and while we do not discuss the details of the particle cascade in this work we

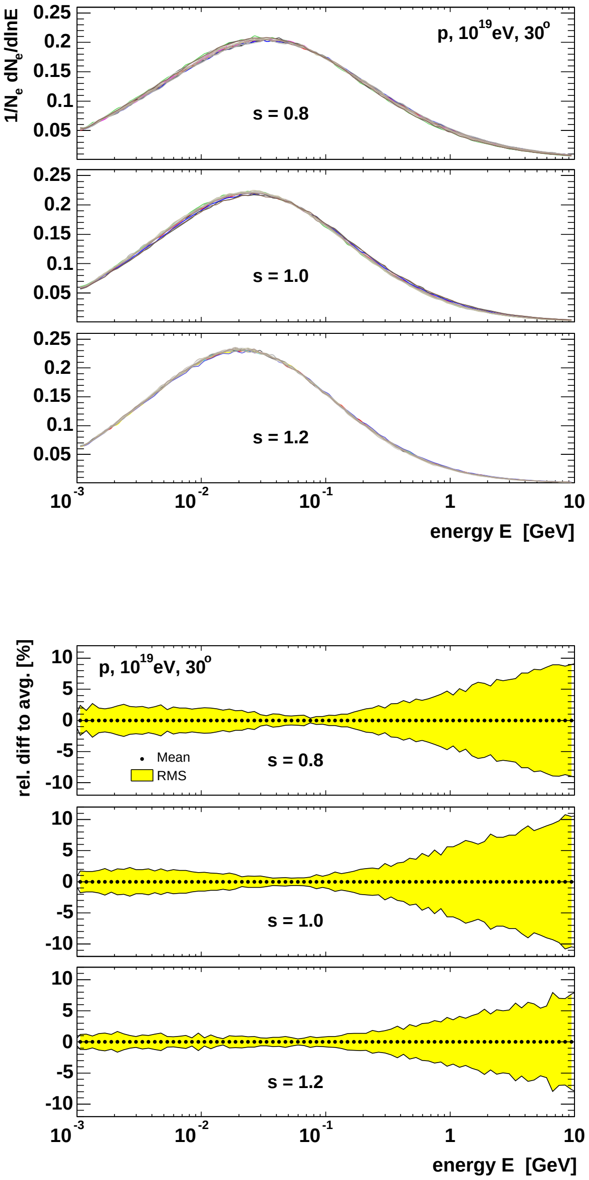

Figure 11: Energy spectra (top) and their fluctuations (RMS) around the mean (bottom) for multiple showers initiated by identical 10^{19} eV protons and sampled at three distinct values of the shower age s. For reference "typical" values of $X_{\max}$ lie around almost 800 g/cm^2 for protons at this energy placing the top plot at roughly 600 g/cm^2 and the bottom one around 1050 g/cm^2 which lies below ground level for the vertical showers we consider in this work. Particles of lower energy penetrate less deeply, for example at 10^{17} eV typical shower maxima lie at 650 g/cm^2. Source: [46]

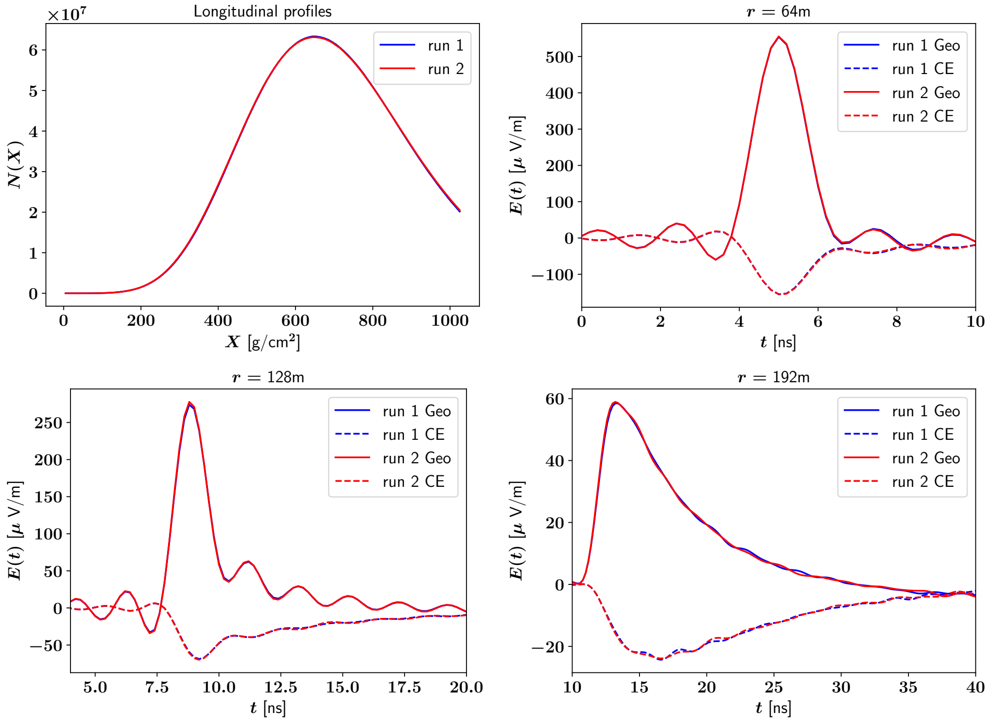

Figure 12: Very good universality for fixed height of first interaction and almost identical longitudinal profiles (shown in the top left). This pair of showers (red, blue) was specifically selected by sampling many longitudinal profiles as such a pairing is quite unlikely. Despite the pulse shape varying greatly between the three lateral positions shown, from within the Cherenkov ring (top right), near it (bottom left) and far outside (bottom right), both Geomagnetic emission (solid lines) and the Charge Excess component (dashed) are virtually identical when comparing the two showers.

must be aware of known limits on universality as they might propagate to the radio signals we are interested in. As an example figure 11 shows energy spectra of electrons and positrons which do fluctuate on the level of 5-10% even when compared at equal shower age.

4.2 RADIO UNIVERSALITY

Based on our macroscopic understanding of coherent radio emission it too should be universal in the same sense the distributions of emitting particles are. Indeed it appears as though showers with identical longitudinal profiles also produce virtually identical observable radio signals as shown in figure 12. But this is only true because we also fixed the propagation geometry of the emitted radio waves. Even assuming perfect universality of the particle cascade equal shower ages s correspond to different depths along the shower axis. For vertical showers the relation between atmospheric depth and geometrical height is simple but nonlinear due to the ambient density itself scaling exponentially with height thus antennas at fixed positions actually observe very different lines of sight to the shower maximum. For an overview of the significance of this geometric scaling refer to figure 13.

While a trivial observation within the context of radio signals it must be emphasised that even a perfectly identical air shower cascade will produce vastly different radio signals in an antenna by propagation effects alone and as these to first order depend on physical distance rather than atmospheric depth we must intrinsically contend with two fundamental scales.

Many reconstruction and analysis techniques implicitly rely on an assumption of universality to navigate this issue in the form of abstracting the entire radio emission as originating from a point source located at the shower maximum [24] and parameterised by its local environment (usually expressed as $X_{\max}$, $N_{\max}$ and the ambient atmospheric density). The radio signal is then assumed to propagate as a coherent spherical wave in the far-field approximation adding a dependence on the line-of-sight distance and the refractive index along said line of sight. While often sufficient within measurement uncertainties (or theoretical uncertainties not directly related to radio emission) this approximation already neglects the evolution of the emitting particle ensemble as the cascade develops. As illustrated in figure 14 even deviations in the one-dimensional longitudinal profile already correlate with visible differences in the radio signal despite these showers being selected for having nearly identical shower maxima. This discrepancy is expected from a theoretical perspective especially for the Cherenkov ring (see equation 5) as this signal amplification arises by emission from all stages of the air shower arriving simultaneously thus leading to constructive interference. Unfortunately said stages are not accessible to measurements but different representations of this concept form the basis of macroscopic predictive models such as MGMR (see section 3.1.3).

At the time of writing there is no truly universal parametrisation let alone a physically motivated analytical model which reproduces measured data within statistical limits. In order to truly represent reality such a model would still require a probabilistic component matching the effects of the remaining fluctuations described in the previous section 4.1 or a detailed study of what its output represents statistically. While aiming to create a novel predictive model concept by studying and expanding upon existing notions of radio universality we gained significant insights into the physics of radio emission from cosmic ray air showers. Our work focuses on the development of the cascade and the interplay between its evolution scale on one hand and the propagation geometry on the other for individual antenna positions. This

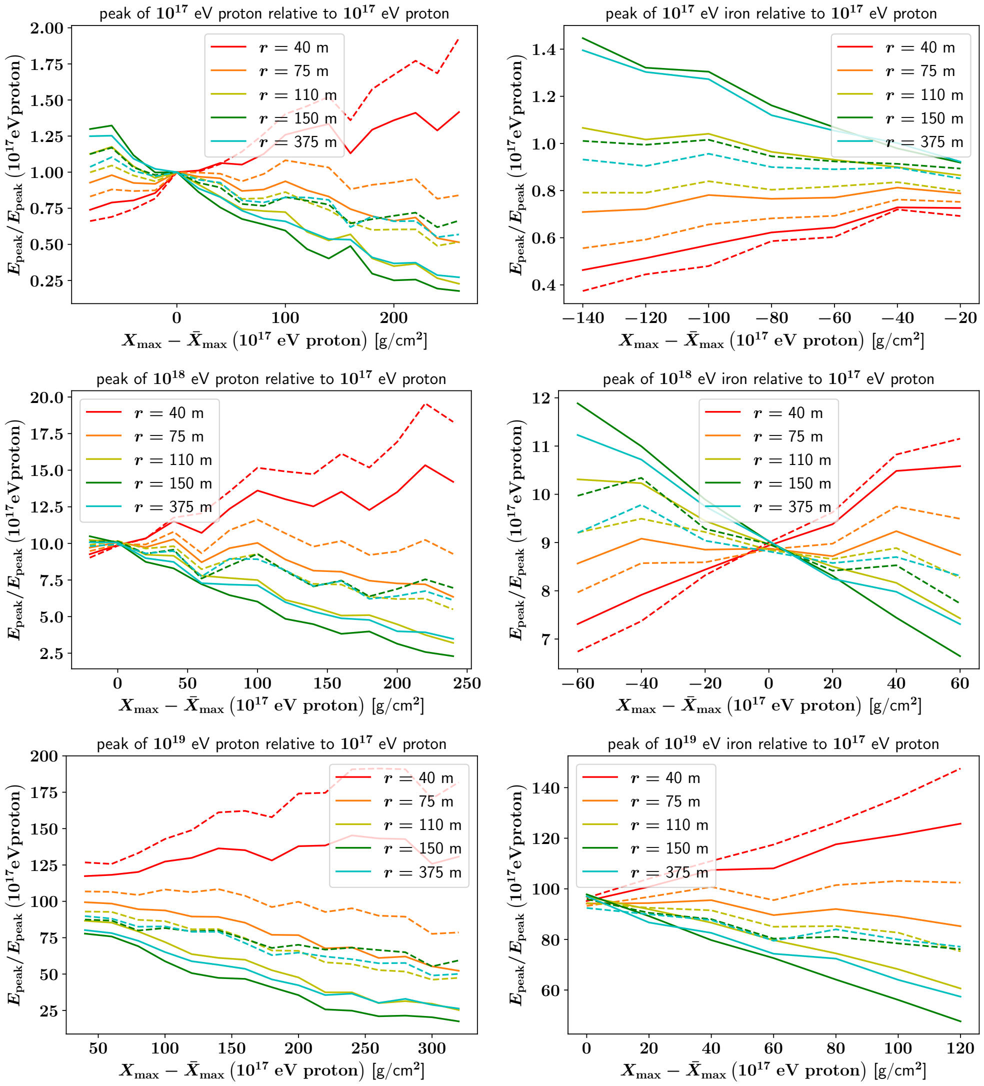

Figure 13: Peak amplitudes of showers initiated by protons (left) and iron nuclei (right) of 10^{17} eV (top), 10^{18} eV (middle) and 10^{19} eV (bottom) relative to a 10^{17} eV proton reference shower. Geomagnetic emission is drawn with solid lines and Charge Excess emission with dashed lines for a wide range of lateral distances. All showers were selected specifically for their depth of shower maximum therefore this spread does not directly translate to observable averages with realistic fluctuations in $X_{\max}$ but our reference shower was selected with an $X_{\max}$ value of 650 g/cm^2 which is close to the mean for 10^{17} eV protons.

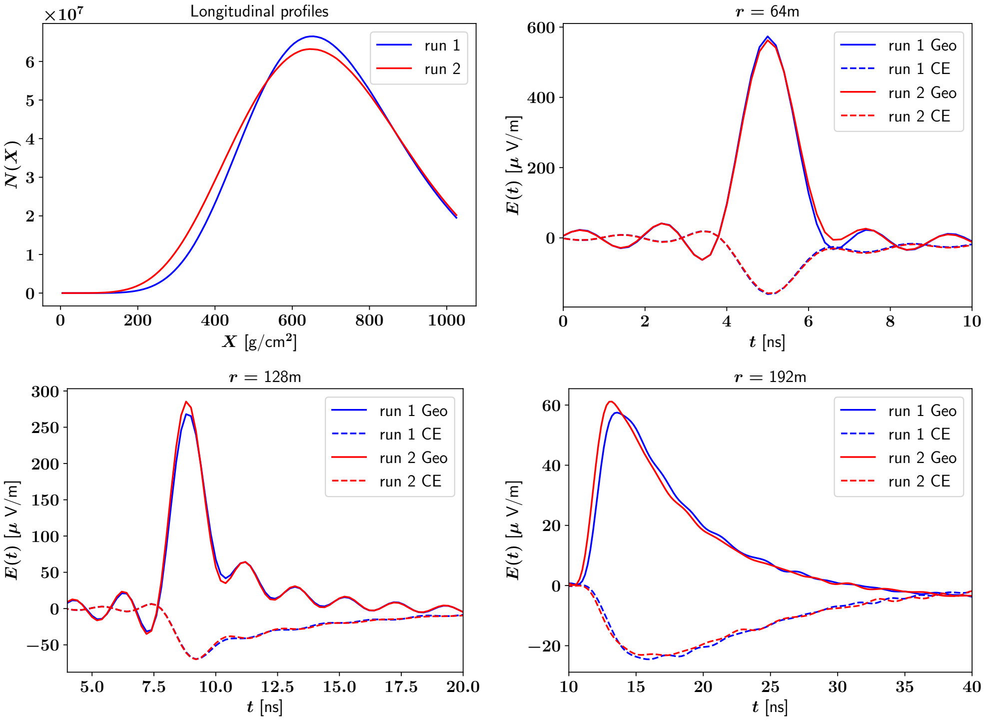

Figure 14: Same layout as figure 12 with the pairing selected for equal $X_{\max}$ but maximising the RMS between the two longitudinal profiles. A simple rescaling with $N_{\max}$ clearly will not match both longitudinal profiles. Differences can be observed in both the peak amplitude and the pulse shape (slope/minima) but the Charge Excess component remains almost identical.

knowledge coupled with the empirical success of our model represents steep progress towards a full universal description in the future.

4.2.1 *Primary Cosmic Ray Energy*

The energy of the primary cosmic ray particle has always been one of the key physical quantities of interest. As the radio signal is almost exclusively produced by the electromagnetic component of the air shower cascade we cannot observe the total primary energy with radio-only measurements. A full *ab initio* radio simulation tool needs to reproduce the highly variable ultra-high-energy quantum processes which determine the fraction of this total energy [48] which ultimately is distributed to electrons and positrons in a manner similar to CORSIKA.

Upon accounting for the total energy contained in the electromagnetic component of the air shower any further transfer into the radio signal appears to be highly universal. According to the well-established shower universality outlined in section 4.1 this energy is highly correlated to the integral or peak height of the longitudinal profile and even most higher-order effects are sufficiently represented by its full shape and these quantities have been directly linked to the total energy emitted at radio frequencies [24].

4.2.2 *Primary Particle Species*

The mass of the primary particle also influences the fraction of its energy which propagates into the electromagnetic component [48]. As for the primary energy determining the mass from measurements of a single secondary component alone usually is not feasible. We will explicitly show comparisons of proton- and iron-induced air showers as the common extreme cases with the expectation of being able to smoothly interpolate between them where quantitative differences arise. On the level of physical observables all nuclei are expected to produce virtually identical radio signals once normalised to the total electromagnetic energy yield of the individual shower.

5

TEMPLATE SYNTHESIS MODEL

Motivated by notions of air shower universality and the success of parametrisations revolving around the shower maximum we set out to substitute the insufficient theoretical source models in classical macroscopic simulation codes with aggregate quantity templates obtained from trusted Monte Carlo simulations. Thus at the cost of the initial template production and processing we hope to obtain a new simulation tool with accuracies comparable to the original template simulations and end-user computation speeds in line with fully analytical models.

The basic premise of our approach is to group the radio emission from many particles into one template. Our core assumption is that these traces can be renormalised to be universal with respect to shower-to-shower fluctuations. Qualitatively the wavelength scale of metres coupled with relativistic beaming should allow such aggregation, however the correct normalisation is not clear. Further the template processing during end-user application ought to be fast to compute and light on memory.

From theoretical considerations we assume an equal amount of radio emission ought to arise from a given distribution of emitting particles at a constant ambient density ρ. In practice identifying and parametrising all relevant secondary properties such as the spatial distribution or individual particle velocities is not feasible, therefore different template processing approaches must be tested and compared pragmatically.

5.1 TEMPLATE CONSTRUCTION

While the cascade evolution is known to depend primarily on the traversed column of matter radio emission instead scales with the geometrical distance between the source and the antenna. Additionally the spatial extent of the emitting cloud of particles determines the coherence of the signal. In conclusion we have two fundamental scales on which to construct our templates and need to evaluate both approaches.

Based on the Monte Carlo cascade calculation in CORSIKA we are able to aggregate microparticle properties into macroscopic quantities. Unfortunately the code is not designed to efficiently follow physical particles as single cohesive entities as expected in classical electrodynamics but instead operates on individual straight “tracks”. Thus our subdivision of the entire shower into template components must use quantities defined for these tracks, the most

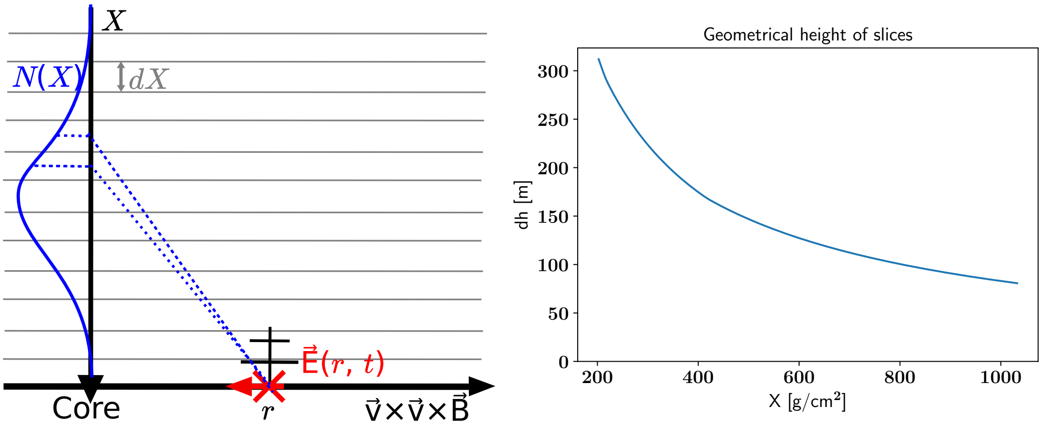

Figure 15: Slice construction in atmospheric depth (left) and conversion of slice thickness to geometrical height dh using the US Standard Atmosphere (right). The air shower processing (blue) is handled internally by CORSIKA+CoREAS, the slice definitions $[X - \delta X/2; X + \delta X/2]$ (grey) are an additional existing input option in released versions and the cascade summary N(X) (blue) as well as the radio time series $\vec{E}(r,t)$ are the primary outputs of the simulation. As the geometrical slice thickness grows exponentially with the height above ground the upper layers of the atmosphere are not displayed in the plot to the right. These do not measurably contribute to the radio signal in most cases due to the very low number of particles.

important one being the positions of their endpoints. We choose to assign tracks to sliced layers of atmospheric (slant) depth X as illustrated in figure 15. Both for technical reasons and due to requiring a sufficient population within each slice for realistic electromagnetic coherence the slices will have a thickness δX of 5 g/cm^2 or larger which corresponds to a minimum geometrical height of roughly 40 m at sea level. The lowest observable frequencies lie around 20 MHz primarily due to atmospheric radio emission [49]. This translates to a maximum wavelength under 15 m meaning we should not be able to spatially resolve individual sources within each slice.

CoREAS assigns radio signals to slices based on the starting point of each track. As long as all eligible slices combined cover the entirety of the air shower cascade the physically observable radio pulse can be recovered by directly summing over all individual slice time series

$$\vec{E}^{\mathrm{phys}}(\vec{r}, t) = \sum_{X=0}^{\infty} \vec{E}^{\mathrm{temp}}(\vec{r}, t, X) \tag{8}$$

where in practice the sum only needs to extend to CORSIKA's lowest observation level which defaults to sea level or roughly 1036g/cm^2 for vertical showers. No propagation time adjustments are necessary as the time reference is tied to the shower front arriving at the antenna location $\vec{r}$ independent of any emitting particles. The slice(s) containing CORSIKA's observation cutoff may contain unreliable data as all particles crossing this horizon are immediately destroyed while in reality they would penetrate the ground and be absorbed relatively quickly. Care should also be taken with antenna positions within the high-energy core of the cascade, while the endpoint formulation remains valid with respect to the track lengths our assumption of macroscopic coherence does not once the distance to the source approaches the measured wavelengths.

In terms of numerics splitting the signal in one physical antenna into hundreds of mutually exclusive template slices does not increase the overall computation cost as each track only contributes to a single slice. However as the number of particles contributing to a slice can easily drop orders of magnitude below the total population it may be necessary to dial back unphysical simplifications despite being adequate for analyses. Our greatest concern falls on numerical "thinning" which groups multiple physical particles into one "effective" aggregate point particle. This can artificially enhance coherence of the radio signal or lead to pronounced spikes in the time domain if not enough numerical particles contribute to a particular slice. Thinning does significantly reduce the computation times of the particle cascade however especially at the energies where radio emission becomes relevant to measurements thus we continue to use it with reduced severity. For large-scale application of our model especially in precision experiments such as LOFAR [3] we recommend the creation of a completely unthinned template simulation library.

5.2 POLARISATION DECOUPLING

As discussed in section 3.1 radio emission should originate from two main emission mechanisms with clearly defined stable polarisation (plus an incoherent contribution which will not be distinguishable from noise). Geomagnetic emission has a uniform direction of polarisation aligned with the Lorentz force $\vec{v} \times \vec{B}$ while Askaryan/Charge Excess emission ought to be radially polarised around the shower axis. This means the two components are orthogonal on the $\vec{v} \times \vec{v} \times \vec{B}$ axis which we can determine *a priori* in our simulations. Therefore we can separate both mechanisms in the raw CoREAS output in perfect agreement with the macroscopic theory and treat them independently where needed, such as only the Geomagnetic component scaling with the ambient magnetic field strength.

An analytic synthesis approach might also have to consider higher order deviations. For example precision measurements verified a time delay between the two radio components manifesting itself as an elliptic polarisation of the total electromagnetic field at the antenna [39]. Since our model works with the full time series data and aims to process both mechanisms separately this effect is innately accounted for provided the separation is achieved. Further there could be an offset of the Charge Excess polarisation symmetry centre with respect to the shower axis, this will require further investigation in the future but no such effects have been observed within the uncertainties of real measurements.

Coherent tilting or rotation of the polarisation pattern or variations during shower development have not been formally excluded and would take significant investments of computation time to study properly. However the electrodynamics underlying the postulation of the base mechanisms do not directly predict such an effect. A minor Charge Excess contribution polarised in the $\vec{v} \times \vec{B}$ direction may also be present due to numerous opportunities for slight asymmetries to arise in the cascade dynamics. As the most prominent large-scale influence is expected from a coherent drift induced by the Geomagnetic field a reliable quantitative study again becomes time-consuming and difficult and will be postponed for future consideration.

5.3 ATMOSPHERIC DEPTH - PARTICLE NUMBER NORMALISATION

For technical reasons fixing the propagation geometry and analytically compensating the cascade evolution happens to be more convenient. In this configuration two or more air showers are simulated with identical physical inputs differing only in their random seeds which in terms of physics primarily determine the depth of the first interaction, the energy transfer into the electromagnetic cascade and the depth of the shower maximum. Template slicing and antenna positions are identical allowing easy permutation of the simulations to increase our sample size of synthesis attempts. For our normalisation we can use the robust longitudinal sampling already done by CORSIKA itself, explicitly choosing to count only electrons and positrons to maintain consistency with CoREAS. However the Monte Carlo cascade cannot efficiently aggregate the lifetimes of physical particles or otherwise easily measure the physical number density of particles. CORSIKA instead defines virtual "detection layers" of fixed atmospheric depth and counts the number of tracks crossing these layers. This is somewhat consistent with the analytical formulation of the longitudinal profile presenting $N(X)$ as an *area density* of particles while our construction of the radio slices must account for all tracks within the *volume* of the slice. Therefore the proper normalisation incurs an additional factor representing the thickness of the slice which will be omitted in the following as we only compare slices of identical shape in this section. An overview of this synthesis configuration is shown in figure 16.

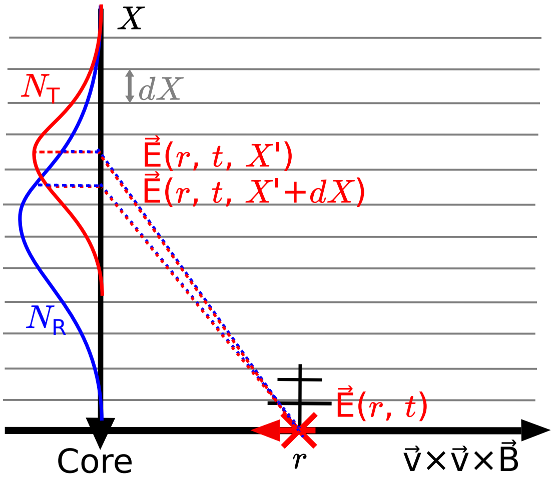

Figure 16: Synthesis layout using slices of fixed atmospheric depth X and constant thickness δX as introduced in figure 15. For each antenna position r on the $\vec{v} \times \vec{v} \times \vec{B}$ axis the two horizontal polarisations are treated independently but rescaled using the longitudinal profiles for the template (t, red) and the target real shower (r, blue) respectively. The lines of sight from each slice to the antenna are identical for all showers thus variations arise solely from the development of the cascade itself.

As this longitudinal profile is a scalar quantity over atmospheric depth we can simply rescale the radio signal originating from each slice individually and sum up the results

$$\vec{E}^{\text{synth}}(\vec{r}, t) = \sum_X \frac{N^{\text{real}}(X)}{N^{\text{temp}}(X)} \vec{E}^{\text{temp}}(\vec{r}, t, X). \tag{9}$$

If the template is equal to the target real shower this expression returns the original Monte Carlo result as per equation 8. As shown in figure 17 this simple approach shows a lot of promise already. Signals at lateral antenna distances far outside the Cherenkov ring in particular are consistently reproduced with near perfect accuracy even though the peak position and pulse shape are emergent properties arising from the sum of many slices. This has particularly promising implications for sparse high-precision detector arrays due to the far greater area illuminated at these distances. Qualitatively the pulse shape of the synthesised time series also appears to closely match that of the real signal even when the peak amplitudes does not. It should be noted the arrival times of the synthesised pulses are *not* fitted to those of the real signal in any way therefore this timing accuracy also represents a direct success of our model.

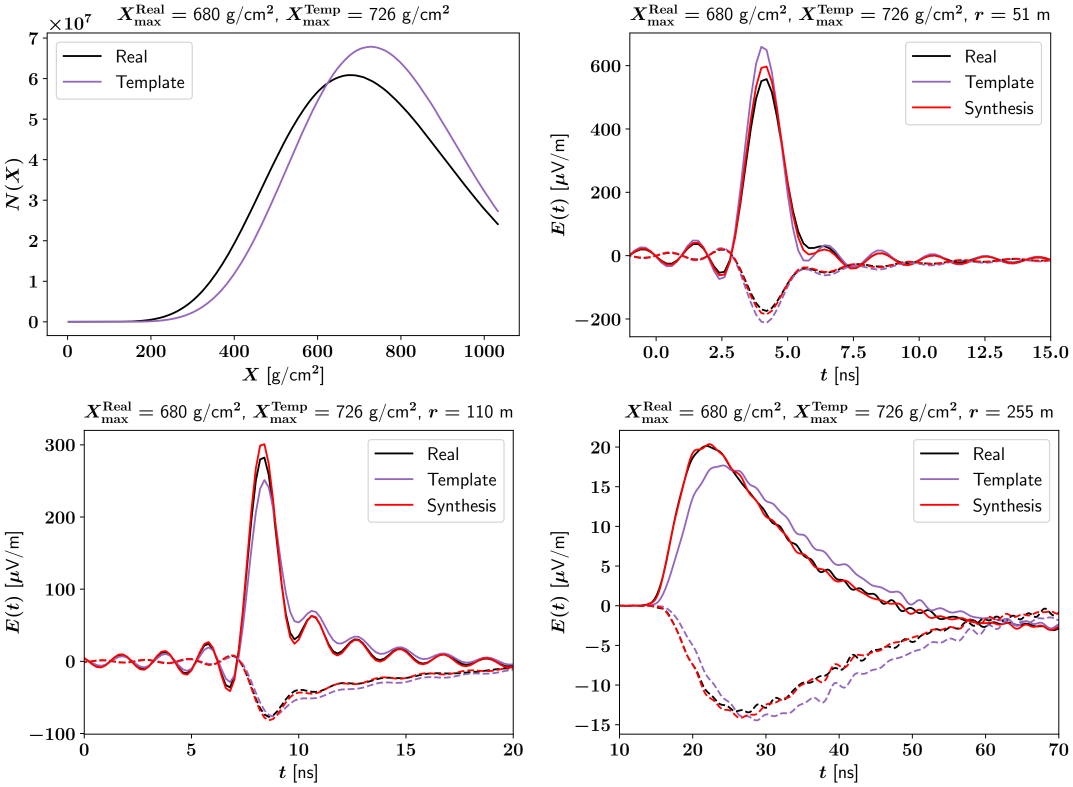

Figure 17: Successful particle number-based synthesis for vertical 10^{17} eV protons and longitudinal profiles used in the synthesis (top left). The observable signals as produced by CoREAS are shown in black for the real signal and in purple for the template, the synthesis result is plotted in red. Solid lines denote the Geomagnetic component, the Charge Excess emission is plotted with dashed lines. While the furthest lateral distance of 255 m (bottom right) is reproduced almost perfectly the inner distance of 51 m (top right) is clearly *under*corrected from the template while the middle distance of 110 m (bottom left) *overshoots* the real signal. Thus an additional refinement with lateral (or longitudinal) extent is necessary to improve accuracy, a global scalar factor representing e.g. the height of the first interaction or the total electromagnetic energy will not suffice.

5.4 LIMITATIONS

Nevertheless some deficiencies quickly become apparent especially when the longitudinal profiles of the template and target are chosen with substantial differences as in figure 18. In particular the peak height at and within the Cherenkov ring does not match the Monte Carlo prediction and this discrepancy has both statistical and systematic contributions. Some statistical fluctuations are expected due to the fundamental limitations of our model such as slice boundaries, imperfect treatment of timings and the universality of various particle properties itself only being valid down to 5-10% according to Lafebre [47]. Some of these may be improved through disproportionate investments of computation time if our approach proves successful overall but especially the last point represents a fundamental limitation to our model. The presence of systematic deviations implies we still have room for improvement before reaching this limit however.

All deviations visibly scale with the differences between the shower maxima of the template and target showers. In figure 19 we show the influence of X_{max} revealing a stable correlation but also significant fluctuations. As this is consistent with the observations in the particle distributions outlined in section 4.1 we assume the cause to lie within the cascade dynamics themselves.

5.5 IMPACT OF SLICE THICKNESS

When grouping microscopic ensembles into macroscopic quantities for use in quasi-analytical models the scale of compression defines a baseline for information loss. In the specific case of slicing particle contributions based on atmospheric depth the total relevant longitudinal range of the air shower cascade amounts to over 10 km and observable wavelengths never exceed 15 m. Thus reaching ideal $\lambda/2$ sampling is barely feasible, and outright impossible for the high-frequency bands: the upper frequency limit of the LOFAR high band lies at 200 MHz or 1.5 m, SKA-low is scheduled to reach 350 MHz or 0.85 m. The latter corresponds to 12000 longitudinal slices for equidistant sampling of 10 km of atmosphere. A further technical limitation which could be overcome with significant investments of computing power applies to any attempt at pre-selecting showers: when used in this capacity as mentioned in section 2.4 CONEX has a maximum resolution of 5 g/cm^2 and systematic accuracy limitations on a comparable scale therefore full Monte Carlo cascade simulations are required instead.

A simple check within the numerically feasible scales has been performed. Given a true subdivision into fine slices one can sum the contributions of multiple slices to achieve a coarser resolution. A mock synthesis using these same data, once with the true slicing and once with artificial coarsening, should di-

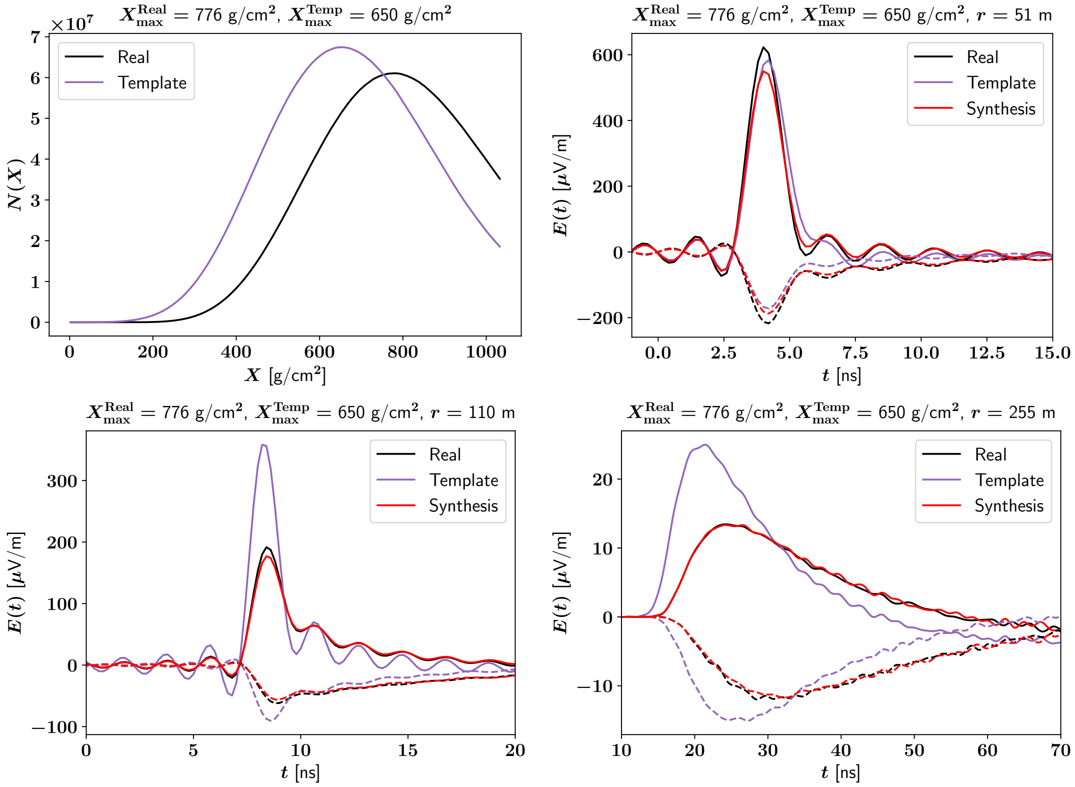

Figure 18: Particle number-based synthesis for vertical 10^{17} eV protons for showers with vastly different $X_{\max}$ following the same layout as figure 17. Despite the shower maxima being over 125 g/cm^2 apart the largest lateral distance once again is matched almost perfectly. Closer inspection of the other distances shows the lack of a consistent hierarchy in the deviations: the Geomagnetic component at the inner distance is not merely undercorrected but rather scaled *away* from the real signal while the middle distance remains overcorrected. As the overshoot occurs in the opposite direction we can assume it to depend on either the ratio of template and target peaks or which $X_{\max}$ lies higher in the atmosphere.

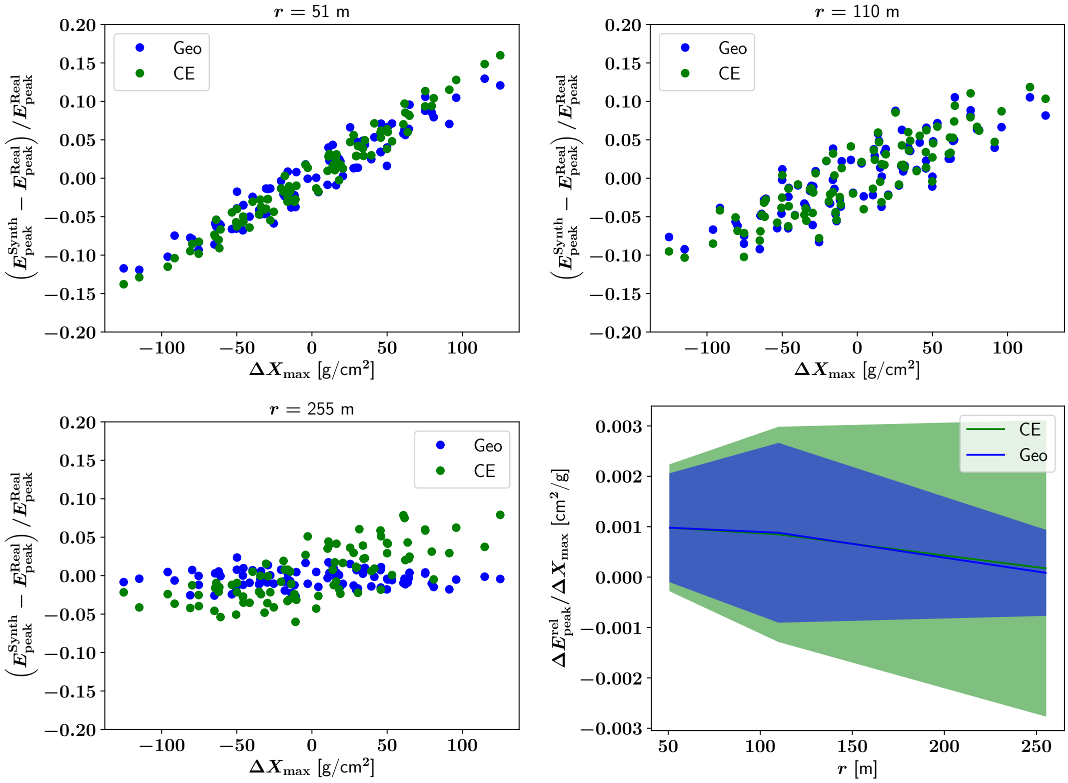

Figure 19: Variation of synthesis residuals with depth of shower maximum. A total of 10 simulations with identical primary particles was used, each once serving as the template to synthesise all other 9 showers. In the first three plots we show the lateral distances used in previous figures. Shown here is the relative difference in peak amplitude between the synthesis result and the real signal over the difference in X_{max} between template and target $\Delta X_{\text{max}} := X_{\text{max}}^{\text{Temp}} - X_{\text{max}}^{\text{Real}}$ for both the Geomagnetic (blue) and Charge Excess (green) components. The final plot shows the mean and 1σ-deviation bands of these peak deviations divided by ΔX_{max}. In this format the reliability of the Geomagnetic synthesis for 255 m truly stands out while comparison with the time series plots in figures 17 and 18 lets us trace the high variance of the Charge Excess peak at this distance to high-frequency wriggles in both the synthesis result and the real signal thus likely a physical incoherent contribution.

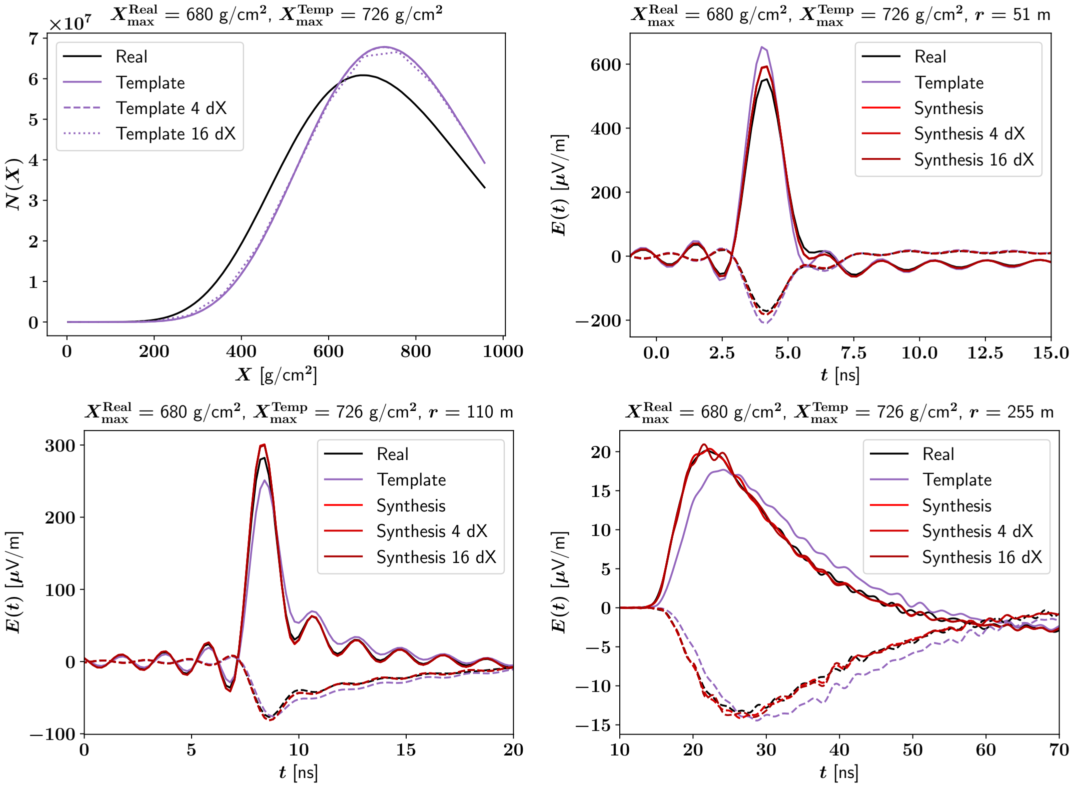

Figure 20: Repeat of figure 17 with the addition of two alternate synthesis curves in darker shades of red: both the longitudinal profile and the CoREAS time series used as the template were compacted using 4 and 16 original slices yielding effective thicknesses δX of 20 g/cm^2 and 80 g/cm^2 respectively. The longitudinal profile was averaged between these values in line with its CORSIKA definition as an area density while the radio time series were summed up which is consistent with the operation of CoREAS.

rectly reveal the effects of lowering the resolution on such macroscopic scales. No significant deviation has been found for an initial resolution of 5 g/cm^2 downsampled to 20 g/cm^2 and only deviations comparable to those already present in our model were observed at 80 g/cm^2 as shown in figure 20. We will continue to use the original template slicing of 5 g/cm^2 going forwards as this empirical observation does not prove an indifference to slice thickness in any aspects of radio emission or shower physics not considered in our simple synthesis model.

5.6 PREVALENCE OF INTERFERENCE

As most existing works focus on the energy fluence or peak height rather than the full time series one important question is to what extent contributions from different slices interfere destructively. For purely constructive interference there is no risk of "magic cancellation" amplifying the significance of small inaccuracies nor of systematic errors in our model negating each other and thus implying a greater accuracy than we actually achieved. Except for the Cherenkov ring these relations also depend on the frequency band under consideration as the peak arrival times from different slices scan a fixed range while the broadening of said peaks depends on the applied filter.

One way of showing the prevalence of interference is to clamp the peak-opposed sign of each individual slice time series to zero before adding them up to recover the "physical" signal as per equation 8. Thus the sum no longer involves any cancellations and the result can be compared to the real signal to gain a qualitative measure of destructive interference. We apply our bandpass filter *before* the clamping in order to include the ringing it causes otherwise the final filtered time series may still include negative values. As shown in figure 21 there is a clear transition where even a fair distance beyond the Cherenkov ring the slices interfere almost purely constructively while for large lateral distances destructive interference can reduce the amplitudes by 50% for both polarisations. It is also apparent that interference affects the shape of the pulse far more than the height of its peak. This is consistent with the theoretical prediction of all partial signals arriving simultaneously at an antenna on the Cherenkov ring as a zero crossing in the timing hierarchy of the longitudinal contributions. The off-peak components of the time series at the Cherenkov ring only originate from edge regions where changes in the value of the Cherenkov angle introduce an additional time shift.

Another way of viewing interference lies in noting the exact time position of the observable time series then plotting the value of each individual slice time series at that specific time over the longitudinal evolution as shown in figure 22. A more complete display would be to scroll through the full time series of each slice but doing so requires animation or interactive media not

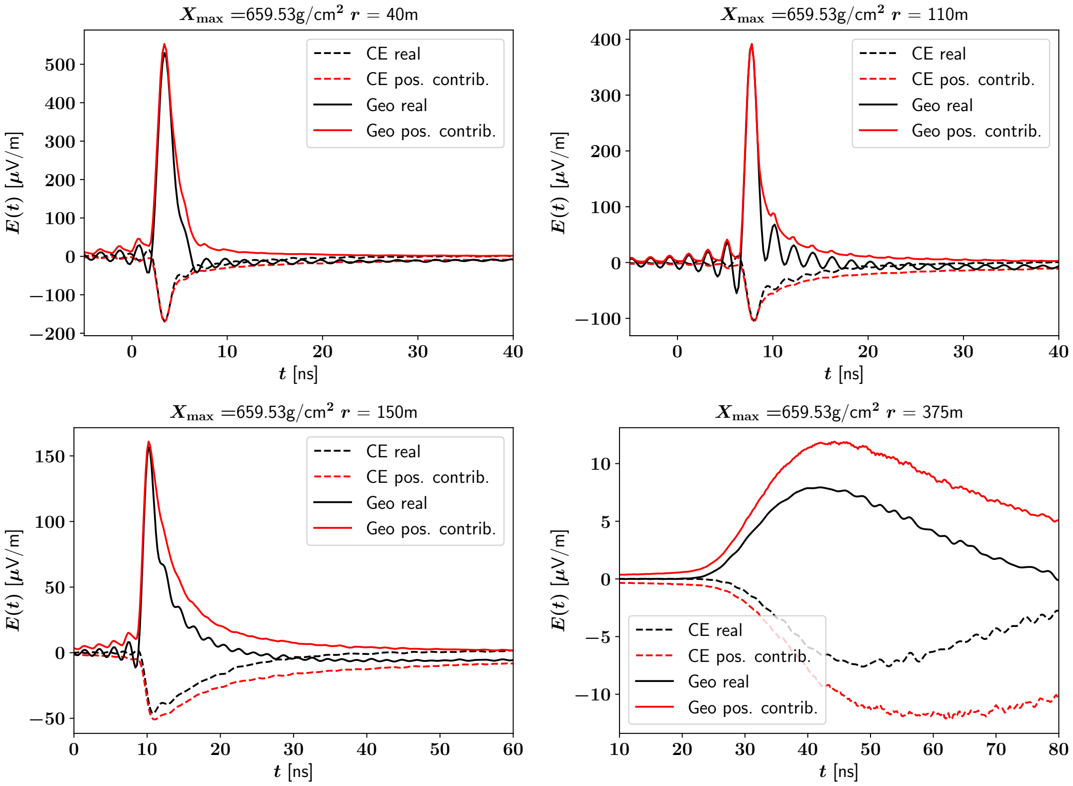

Figure 21: Physically real (black) and interference-omitted (red) radio time series for both Geomagnetic (solid line) and Charge Excess (dashed) emission for different lateral distances.

properly reproducible in print. Most of the knowledge gleaned from such animations can be identified in the contributions to the physical peak however. In general each slice time series very closely resembles a delta function (or the bandpass filter response to it) and progressing in the longitudinal direction scans a continuous evolution of arrival times and amplitudes. The longitudinal range contributing to the physical peak largely precedes the shower maximum but tends to be much broader than common X_{max} resolutions meaning the common approximation as a point source in many analyses is inaccurate. One major reason it empirically functions regardless lies in the small amount of destructive interference coupled with the fact most propagation effects merely shift signals to slightly different times or lateral distances. Because these analyses tend to focus on measuring the energy fluence or even just the total radiation energy in the entire radio footprint these shifts do not affect the integral quantities leaving only a smooth correlation.

The exact curvature and especially the zero crossings are defined almost exclusively by our bandpass filter. Progressing in the longitudinal direction effectively scans the time-domain response pattern of the filter, modulated by the strength of the actual radio signal. The significantly shorter longitudinal range contributing to the largest lateral distances is *not* specific to our frequency band however and a comparison with the time series in figure 21 shows much of the pulse energy is distributed from the peak to a broader bulk.

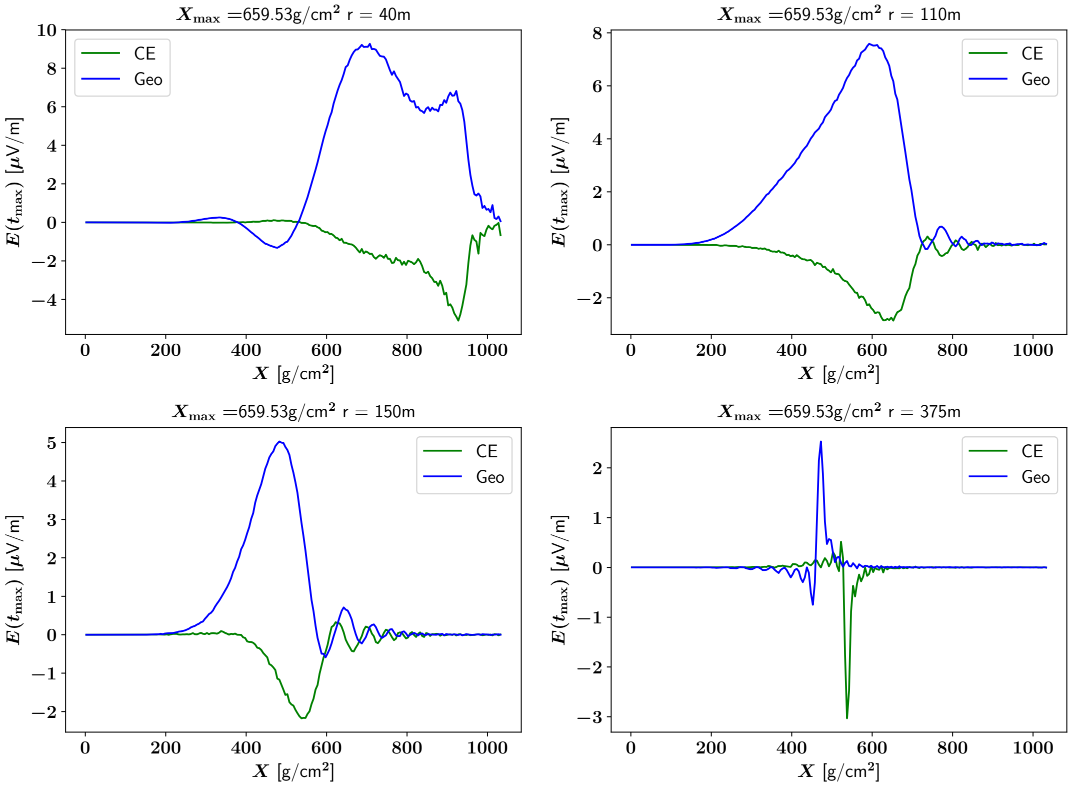

Figure 22: Longitudinal contribution from each slice to the time bin of the physically real peak for both Geomagnetic (blue) and Charge Excess (green) emission over the same lateral distances shown in figure 21. For reference the shower maximum of this particular shower viewed under the Cherenkov angle equates to a lateral distance of 80 m at ground level. We have no simple explanation for the second peak in the contributions at 40 m but assume it arises from a combination of cascade physics and a steady increase in the value of the Cherenkov angle beaming more radio emission from the very late stages of the shower to larger lateral distances.

6

LONGITUDINAL EVOLUTION ALTERNATIVES

The simplest addition to our model is to modify the scalar factor applied to each individual slice time series. There are several potential motivations to do so, for example the common longitudinal profile only counts the raw number of particles but not their momenta which can significantly influence coherence. Following from the instantaneous velocity distribution the spatial distribution of the particles also changes as the air shower develops. Both of these phenomena depend on the evolution of the cascade moreso than the absolute geometry of the system which is consistent with our observed discrepancies.

6.1 RELATIVE EVOLUTION PARAMETRISATION

Again following Lafebre et al. [47] proper universality is only achieved if all shower maxima are overlaid in the parametrisation. While the particle cascade appears to have only one defining scale radio emission also is subject to propagation physics. In our simple test the absolute geometry for the template and target were identical. This convenience will need to be abandoned eventually but it is important to know how relevant the absolute geometry and the propagation of the radio signal actually are compared to the internal developmental scale of the particle cascade.

Based on the definition of the relative evolution stage in equation 6 we should compare slices at the same distance (in atmospheric depth) from the respective shower maxima. The radiation length can be neglected as it is assumed to be constant therefore we define

$$\Delta X = X - X_{\text{max}} \tag{10}$$

as our new longitudinal parameter. Anticipating the full angular emission spectrum of our Monte Carlo radio sources is futile. Since the effective refractive index does not change significantly over typical differences in X_{max} we instead compare matching viewing angles ε between the shower axis $\vec{v}$ and the line of sight $\vec{l}$ to the antenna drawn from the centre of the slice. In addition we need to choose a reference distance along this line of sight, the most promising options being the projection down to ground level which most closely matches physically observable configurations on one hand and equal geometric distances to account for amplitude scaling due to flux conservation on the other. Due to a lack of prior works dealing with unphysical

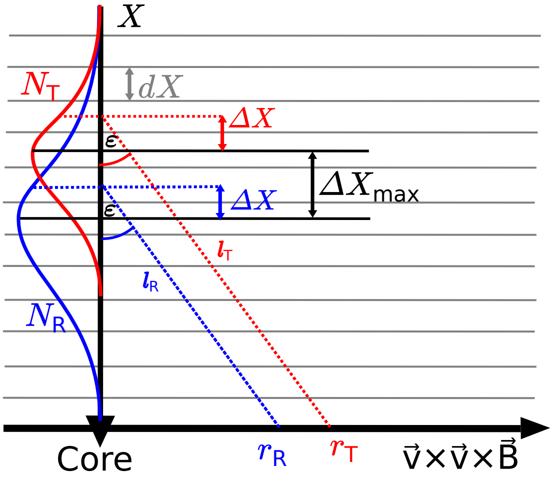

Figure 23: Synthesis layout using slices of equal distance ΔX to the respective depth of shower maximum X_{max}. In this example the antenna position of the template r_T is projected down to ground level and therefore at a different lateral distance than the target real position r_R also resulting in different line-of-sight lengths for the template l_T and target l_R. Templates are instead matched by the viewing angle ϵ between the shower axis and the antenna (which is also equal in the atmospheric depth slicing as per figure 16). In this configuration the template depends not only on its own shower cascade but also the target or more specifically the *difference* in the depths of the shower maxima ΔX_{max}

antenna positions and potential pitfalls in treating the atmosphere below sea level when shifting downwards we predominantly favour the former. For a schematic of this template construction refer to figure 23.

As our template configuration now depends on the cascade realisation we can no longer easily exchange templates and target showers. The minimum computation investment for a template, while still equal to a “real” shower, no longer produces a physically observable signal at ground level. Worse still the projected lateral distances (or heights if comparing equal line-of-sight lengths) must be tailored to a specific value of ΔX_{max} i.e. a specific *pair* of template:target showers which severely limits our statistics.

6.2 LINE-OF-SIGHT PROPAGATION - FAR FIELD

When using template antenna positions projected to ground level the length of the line of sight from the source differs to the equivalent in the real shower. This needs to be acocunted for in both the arrival times as the pulses propagate as well as their amplitude. Based on conservation of (energy) flux from an ideal point source we expect the field strength of the radio signal to scale inversely with the distance to the source. The pulse front should still propagate with the speed of light in the atmosphere, but this speed may vary as the refrective index is not homogeneous.

In order to apply our modifications it is important to note the adjustments CoREAS already makes. The global time reference $t = 0$ denotes the moment the theoretical shower front impacts the ground. For our purposes this will always be a constant as all detector positions are located at sea level. What does change however is the *internal* beginning of the timeline which is defined by the depth of first interaction X_1. Thus the earliest physical time referenced by CORSIKA and CoREAS is $-h_1(X_1)/c$. From this reference point CoREAS draws straight lines of sight $\vec{l}$ to each given antenna position and automatically adjusts the numerical evaluation window (but *not* the physical time axis!) around the moment information could have traversed this line of sight i.e. l/c. Due to the finite length of the simulated time series this may cause contributions from (physically unimportant) regions of the cascade to be omitted entirely or conversely some unphysical antenna locations to register observable signal contributions at times outside the time window as defined at ground level. The system generally is designed to always capture the bulk of the signal at the explicitly given antenna location therefore we can project them as needed.

As long as the arrival time adjustments remain small compared to the total length of the sampled radio time series we can apply them to the electric field in the frequency domain as a phase gradient $\Delta\phi$. In order to add up contributions from multiple slices we need to line up discrete numerical time bins therefore direct shifting in the time domain would cause frequent rounding errors on the scale of our time resolution of 0.2 ns. Separate bookkeeping is required for larger differences in arrival times as the absolute time scale remains defined solely by the template shower.

The pulse amplitudes also require adjustment. In accordance with our approximating individual slices as point sources on the shower axis the electric field should scale with the inverse of the geometric line of sight length l. Neglecting the minor curvature induced by inhomogeneity of the refractive index (which CoREAS does not account for either) our highly symmetric shower geometry allows us to substitute ratios of these distances with the corresponding heights or lateral distances as per the intercept theorem. It is important to note that while CoREAS internally calculates only the far-field contribution in the full Liénard-Wiechert potentials this approximation is made in regards to the length of *individual particle tracks* not the distances between them. Therefore the coherent sum could in principle display a different scaling behaviour.

6.3 DIPOLE EMISSION NORMALISATION

In addition we now need to reexamine our understanding of radio emission physics. Comparing equal absolute geometries also has the side effect of comparing equal ambient densities at the source location. While any effects on the longitudinal development of the cascade should be accounted for by parametrising our templates in atmospheric depth instead of geometrical size the ambient density also affects the lateral spread and velocity distributions of the particles which are crucial for coherent radio emission. To first order radio emission from electrons and positrons scales with their mean free path length which is inversely proportional to the ambient density. Additionally we must pay close attention to the definition of the particle count we use to define our longitudinal profile: CORSIKA counts *tracks* which cross a virtual observation plane of fixed atmospheric depth. Physically this represents an *area density* of particles rather than an absolute count. The slices meanwhile are defined as observation *volumes* with infinite lateral extent but finite thickness. Omitting global constants and units this amounts to an additional factor inversely proportional to the atmospheric density representing the geometric thickness of each slice.

Treating slice geometry and particle dynamics independently may not be the most accurate description. Based on classical electrodynamics we also computed various aggregate quantities within each slice more closely aligned with our theoretical understanding of "sources". Among these the most promising amounts to a charge-weighted sum over the geometric track vectors

$$\vec{D}(X) = \sum_{i}^{N_{\text{particles}}} q_i \cdot \left(\vec{d}_{\text{end}}(i) - \vec{d}_{\text{start}}(i)\right) \tag{11}$$

where q_i is the particle's charge, $\vec{d}$ the spatial coordinates of the track's endpoints and the sum extends over all particles starting within a slice of constant thickness δX at atmospheric depths from $X - \delta X$ to $X + \delta X$. The slice borders were chosen to be consistent with the slicing of the radio signal. Physically the vector-valued quantity $\vec{D}$ represents an effective dipole moment within the slice of which the second component (aligned with the $\vec{v} \times \vec{B}$ axis) should correlate with Geomagnetic emission and the third component (along the shower axis $\vec{v}$) produces the Charge Excess signal. As this dipole moment was averaged over the full volume of the slice no additional normalisation is required when comparing on the scale of atmospheric depth.

6.4 RELATIVE EVOLUTION - DIPOLE SYNTHESIS

Combining the template construction from section 6.1, the line-of-sight corrections from section 6.2 and the effective dipole norm from the previous section we assemble a refined synthesis model. As the comparison between figures 17 and 24 shows we can achieve slightly better accuracy for the peak amplitude and pulse shape despite tackling a less restrictive idealised scenario.

Unfortunately this approach immediately encounters several issues and limitations. While an improvement within the Cherenkov ring it performs noticeably **worse** and less consistently for larger lateral distances. Further some failures produce clearly unphysical pulse shapes with no obvious numerical origin such as division by zero or similar divergences. Additionally the accessible template space is limited by our observation level which becomes a major drawback if the template lies lower down in the atmosphere relative to the real shower. Essentially this issue is identical to the "clipping" [24] which becomes a concern for real measurements at higher elevations only in our configuration the template is cut off at sea level and our virtual antennas would need to lie deep below where the atmosphere model becomes untrustworthy.

A fairer comparison with our previous model would involve disentangling the various adjustments but applying the effective dipole norm to templates of equal absolute atmospheric depth would double-correct for emission characteristics. Meanwhile forcing a disjunction between the absolute depth slicing and the propagation geometry requires turning to the even more complex case of differing zenith angles or the use of unphysical settings such as artificially increasing the density of the atmosphere. We investigated several potential angles but were unable to derive meaningful knowledge as the programming restrictions of CORSIKA and lack of external verification options left severe doubts over the validity of the simulations themselves let alone the assumptions underpinning our template model being fulfilled.

In examining the line-of-sight scaling as a necessary component addition over the original model we uncovered a general limitation regarding our approximation of individual slices as ideal point sources outlined in the following section 6.5. As the comparison of slices at vastly different atmospheric depths (and thus geometrical distances) greatly amplifies any errors in this rescaling we chose to abandon the approach of relative evolution synthesis entirely. Future works may attempt to parametrise the nonlinear scaling of the radio amplitude with the line-of-sight distance in the pursuit of theoretical understanding but the implied reliance on a poorly understood paramteric fit *in addition* to the technical issues lead us to pursue different avenues instead.

6.5 POINT SOURCE APPROXIMATION - NEAR-FIELD EFFECTS

For a given antenna position the longitudinal synthesis approach assigns one radio time series to a macroscopic slice of finite vertical thickness but infinite lateral extent. Physically this is equivalent to approximating the emitting region as a point source with an arbitrary directional emission pattern. We

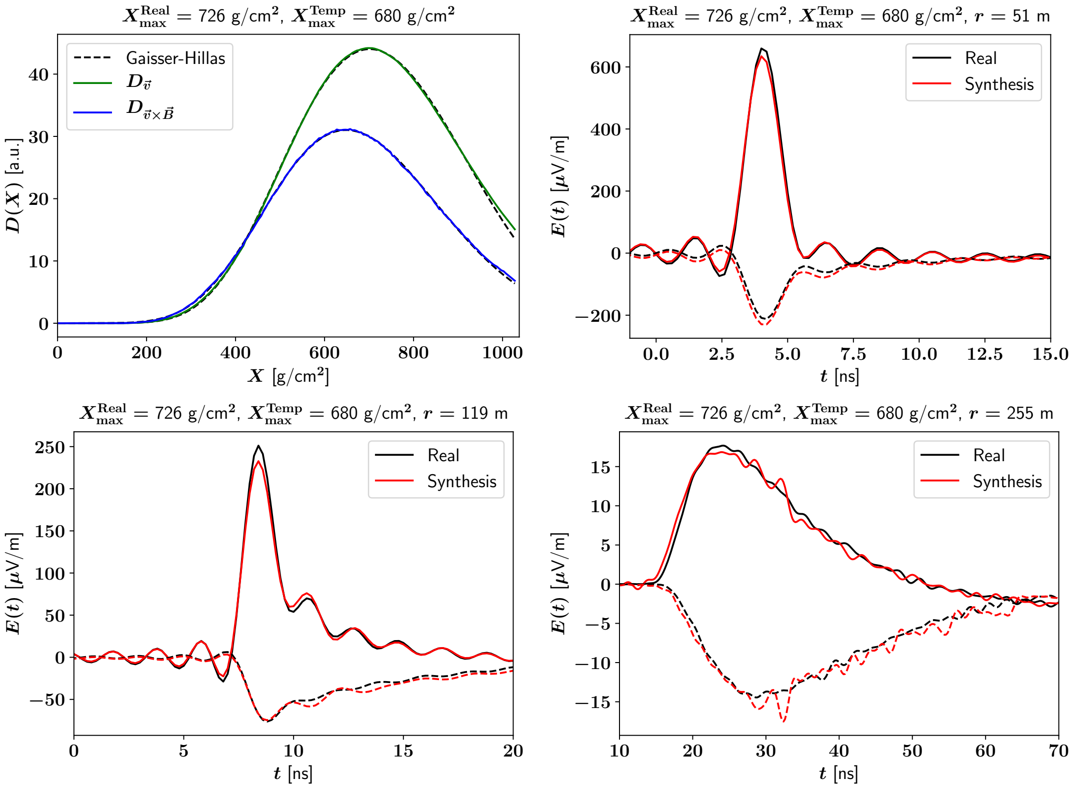

Figure 24: Synthesis of the same shower pairing used in figure 17 matching slices by distance from the depth of shower maximum ΔX and using the slice-averaged dipole moment $\vec{D}$ for rescaling. The $\vec{v} \times \vec{B}$ component of $\vec{D}$ ($\hat{y}$, blue) was used to synthesise the Geomagnetic component (solid line), the $\vec{v}$ component of $\vec{D}$ ($\hat{z}$, green) for the Charge Excess emission (dashed line). Both components of $\vec{D}$ can be fitted very accurately using the same Gaisser-Hillas formula defined in equation 4 for the longitudinal profile of cascade particles. The odd "spikes" in both polarisations at 255 m are the product of incorrect interference between multiple slices rather than an obvious scaling error in a single slice. Similar features arise at different times for other pairs of showers including at the peak of the time series making this synthesis model unreliable regarding the shape of the produced pulses.

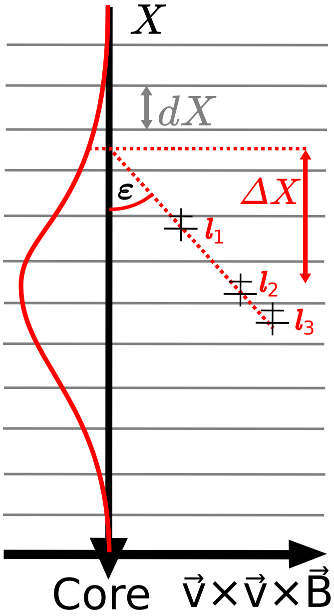

Figure 25: Simulation layout for our line-of-sight test. A few slices ΔX_i at fixed distances from X_{max} are selected to save computation time. For each slice several lines of sight with viewing angles ε are drawn in the $\vec{v} \times \vec{v} \times \vec{B}$ direction. Each line of sight is then sampled at set distances l_j from the slice centre positioning most antennas at unreachable locations in the sky or underground.

can test this assumption by comparing the radio signal from a given slice at different points along a fixed line of sight using physically unreachable antenna positions as shown in figure 25.

In total we compare five emitting slices ΔX_i located at a fixed offset from the (known, true MC) X_{max}, these being

$$\Delta X_i = [-200; -100; 0; 100; 200] \, \frac{\text{g}}{\text{cm}^2} \tag{12}$$

which roughly covers the FWHM range of the particle number distribution. Offsets to even later stages of shower development are not universally feasible as some values of X_{max} may shift the slices themselves below sea level. The range of antenna distances would also be constrained due to hardcoded limitations in CORSIKA's atmosphere model: it is only extrapolated down to 5 km below sea level.

From each slice centre (in terms of atmospheric depth) a line of sight $\vec{l}$ is drawn for each of the following angles in the plane spanned by the $\vec{v}$ (CoREAS $-\hat{z}$) and $-\vec{v} \times \vec{v} \times \vec{B}$ (CoREAS $+\hat{x}$) axes

$$\varepsilon = [0; 0.1; 1; 2; 3; 5; 10]^\circ \tag{13}$$

where the Cherenkov angles from X_{max} to ground level range between 1.18° and 1.37°. Deliberate scaling around the Cherenkov angles for each slice would also be possible but the absolute angular scale more closely represents the geometrical aspects of interest.

The radio signal along each line of sight is then sampled at fixed points l every 500 m. For a perfect point source we expect the pulse to propagate for a duration of

$$\Delta t = l \cdot \frac{n_{\text{eff}}(l)}{c} \tag{14}$$

where n_{eff} is the effective refrective index calculated using the same estimate as implemented in CoREAS itself.

The electric field amplitude of a point source meanwhile drops with increasing distance as

$$E(l) = E(0) \cdot \frac{1}{l}. \tag{15}$$

Applying these scalings in reverse yields predictions of the electric field at the source location which can be directly compared to each other.

As shown in figures 26 and 27 larger distances match almost perfectly but at distances shorter than 3 km our approximation begins to break down. Notably only the 3-4 closest curves (matching distances between 500 m and 2 km) are clearly distinguishable by eye after which the entire pulse shape stabilises validating our simple scaling approach for much of the air shower. Timing discrepancies are small (< 1 ns) and could almost exclusively be caused by the interaction between our discrete time sampling at 0.2 ns and the continuous analytical application of equation 14 in our correction.

For small viewing angles a global amplitude scaling might be sufficient while at larger angles the pulse shape also changes with distance to the source. The breakdown point of ≈ 3 km corresponds to an atmospheric depth of ≈ 700 g/cm^2 when raised vertically off the ground at sea level.

Additionally the ratio of short- to long-distance peak amplitudes shown in figure 29 shows an inversion between 1° and 2° which may relate to the Cherenkov angle but does not appear to orient directly around its value.

In conclusion our timing model works extremely reliably even close to the source. The peak amplitude deviates from simple $1/l$ scaling at distances $l < 3$ km with a nontrivial angular dependence. As the amplitude scaling shows a high degree of universality between different showers as well as in the longitudinal direction it may be possible to parametrise it. A better approach for future investigations would be to identify the physical origin of this scaling within the cascade.

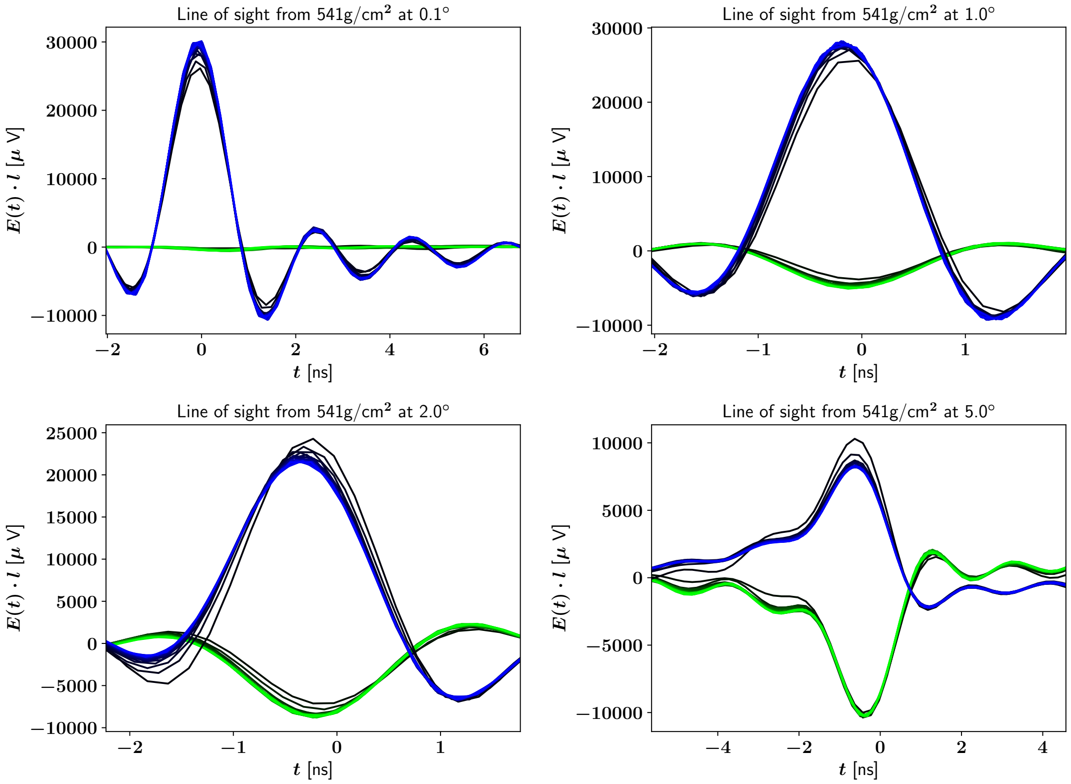

Figure 26: Example LoS-rescaling on the time series level for a slice 200 g/cm^2 above X_{max} and multiple viewing angles. Geomagnetic emission is coloured blue, Charge Excess emission green with both starting at maximum brightness at the largest distance $l = 10$ km and fading to black as they approach the source. The time axis is back-projected not only along the line of sight as per equation 14 but also vertically down to ground level (with velocity c) introducing a global shift for better readability. The electric field is only scaled with the LoS length l.

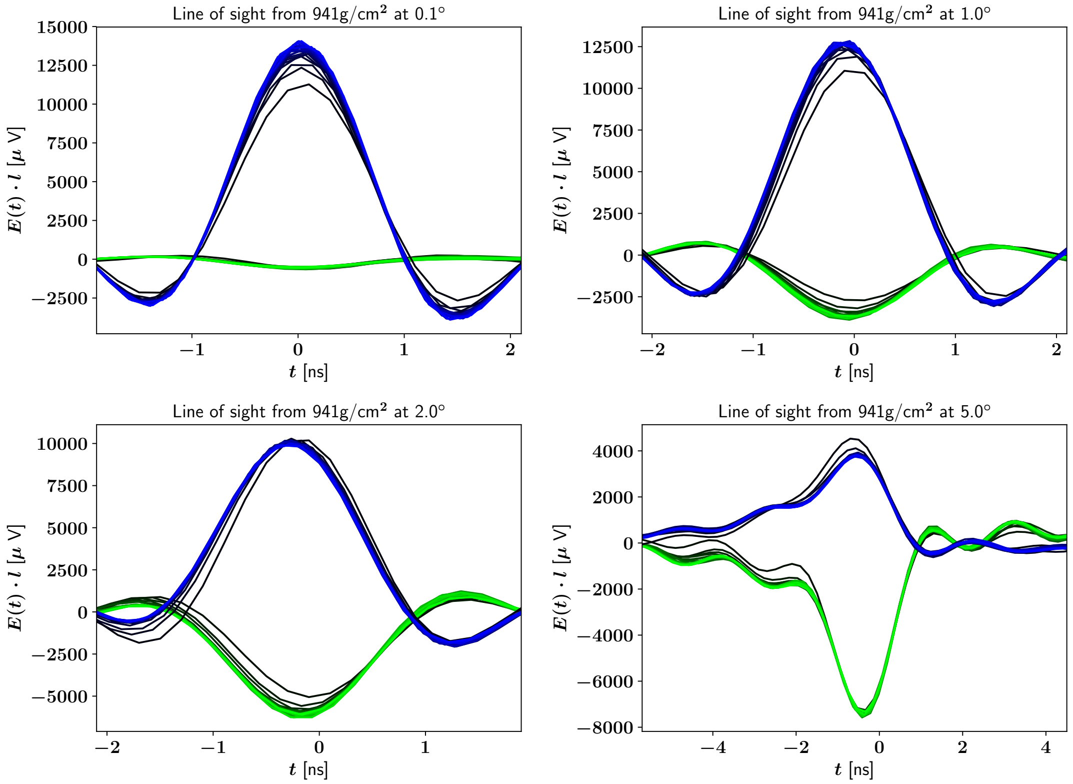

Figure 27: Same as figure 26 but for a slice 200 g/cm^2 below X_{max}. In this case the antennas mostly lie at unphysical positions below sea level.

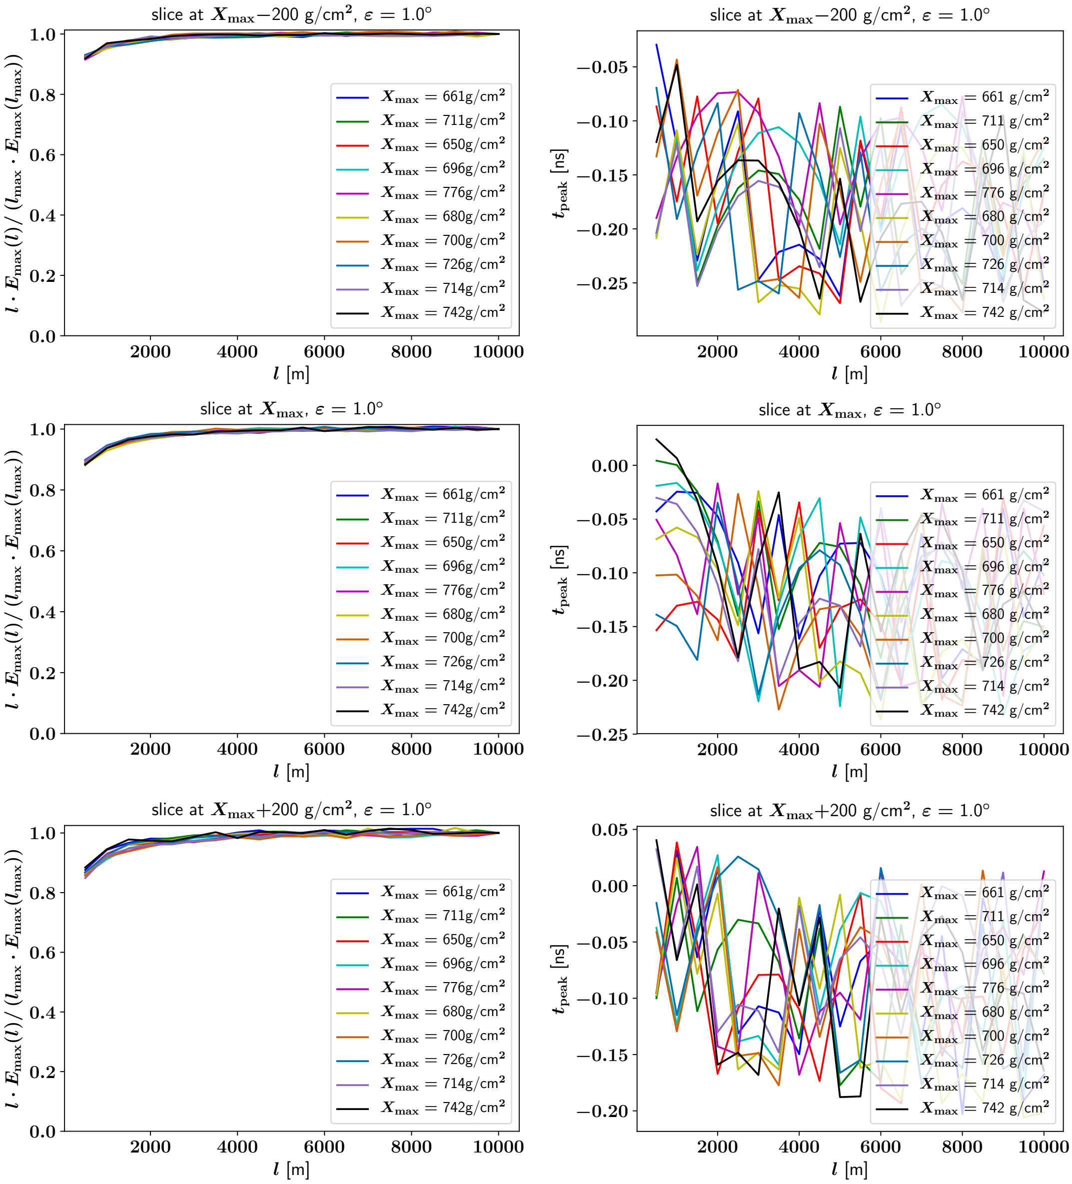

Figure 28: Comparison of peak amplitudes and arrival times across multiple showers for a single viewing angle of $\varepsilon = 1°$ and slices 200 g/cm^2 above and below as well as at $X_{\max}$. The peak amplitudes are divided by the value at the largest distance for each shower individually therefore perfect overlap here is expected while continued overlap along the line of sight shows universal scaling behaviour across all showers. The arrival times were adjusted according to equation 14 and then vertically down-projected to ground level with velocity c to overlap slices from different showers. This reaches our numerical time resolution of 0.2 ns which limits the determination of the peak position.

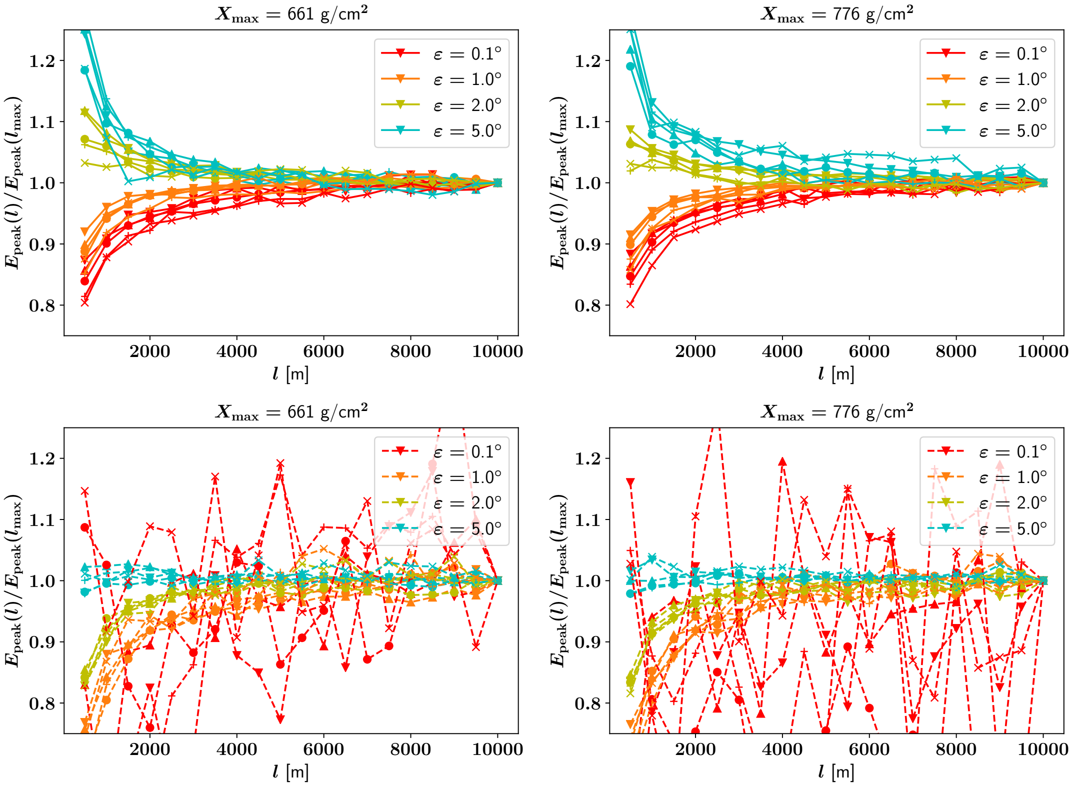

Figure 29: Plot of all time series peak amplitudes relative to the furthest sample along their respective line of sight. The colours denote different viewing angles (red: 0.1°, orange: 1°, yellow: 2°, cyan: 5°) and the markers represent different source slice offsets (triangle down: $X_{\text{max}}-$ 200 g/cm^2, triangle up: $X_{\text{max}}-$ 100 g/cm^2, circle: X_{max}, plus: $X_{\text{max}}+$ 100 g/cm^2, cross: $X_{\text{max}}+$ 200 g/cm^2). Note the largest viewing angles are far beyond realistic observation scales and despite filtering and upsampling the numerical determination of the peak position is somewhat unstable. The top two plots show the Geomagnetic component, the bottom two Charge Excess emission. The left column shows the same shower used to produce figures 26 and 27, the right column shows an additional run with a far greater X_{max}. The high fluctuations for $\varepsilon = 0.1°$ hint at incoherence or very low absolute amplitudes. As each line of sight is normalised individually their overlap does not in any way imply universality as previously discussed in section 4.2.

6.6 RADIO CONTRIBUTION EVOLUTION

The construction of our line-of-sight test can also be used for visualisation purposes. By positioning antennas at the same distance and under the same viewing angle for every slice we can essentially sample the radio emission strength of the cascade in the longitudinal direction. This configuration seems the most promising for further studies of universality especially as modifications of the CoREAS source code can be used to group and slice particles by multiple additional properties.

For our purposes the main benefit lies in directly comparing this "radio longitudinal profile" to the normalisation curves in our synthesis model. Direct equality is not *necessary*, as shown in section 5.6 and implied by the results of the relative-evolution-sliced synthesis variant interference between multiple slices can cancel out inequalities via longitudinal correlations within the cascade. Finding a physically motivated description which matches all radio profiles directly would certainly be *sufficient* for a reliable synthesis model however.

In figure 30 we show the longitudinal profiles of several typical viewing angles which appear to be quite similar in shape but differ between the two emission components. We also included two particle longitudinal profiles with N/ρ representing a properly normalised volume density of electrons and positrons and seemingly aligning with the Charge Excess profiles though they do also fluctuate slightly in their overall shape. With an additional factor of $1/\rho$ the profile aligns with the Geomagnetic component instead which may represent the mean free path length modulating the geomagnetically induced drift current as assumed in the classical models in section 3.1.

6.7 EMITTING PARTICLE ENERGY

One aspect we have not been able to study in detail but which may prove important in the future is the kinetic energy distribution of the electrons and positrons producing the radio signal. To begin with according to figure 11 the contributions from particles above Lorentz factors of $\gamma \approx 1000$ may be subject to shower-to-shower fluctuations on the order of 5-10% which fundamentally limit the theoretically achievable precision of any template-based synthesis approach. As illustrated in figure 31 these high-energy particles do in fact contribute significantly to the total radio emission but lower logarithmic bins in γ contribute roughly equal amounts meaning their greater universality should dampen the effects of fluctuations.

We also determined this observation holds on the level of individual slices as exemplified by their contribution to the physical peak in figure 32. On closer inspection some longitudinal variations are visible thus it might be

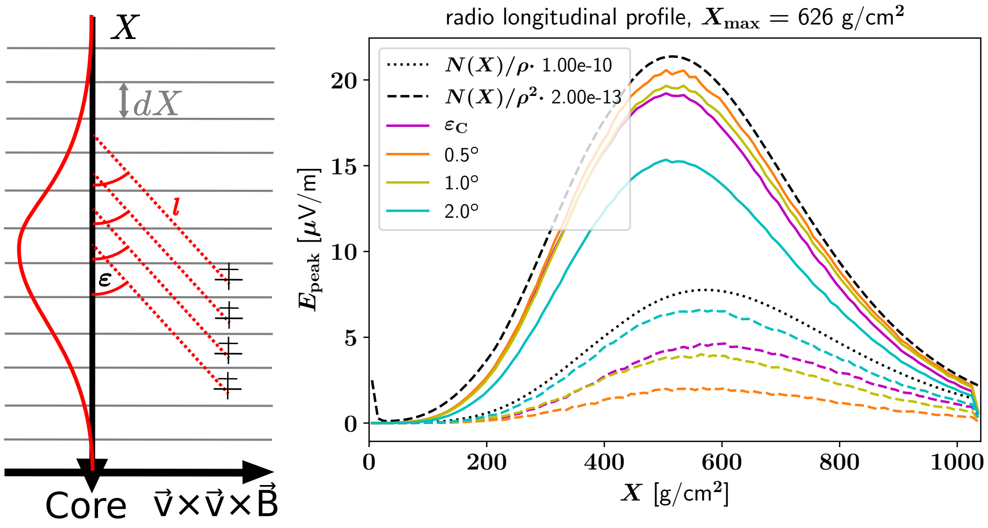

Figure 30: Left: Schematic for sampling longitudinal radio profiles based on figure 25 but now including all slices with only one line-of-sight distance l each. Right: Resulting radio profiles for multiple viewing angles ε in the same shower using the peak amplitude. Geomagnetic emission is plotted with solid lines and Charge Excess emission with dashed lines both at their true physical magnitude. The small-scale fluctuations visible in all curves predominantly are of numerical origin as the discrete times sampled by CoREAS may sometimes coincide directly with the physical peak while at other times the peak lies in the middle between two bins. Contributions from incoherent signals and numerical noise are also possible. We also plotted the electron+positron longitudinal profile $N(X)$ (black) rescaled with the ambient density $\rho(X)$ and an arbitrary scalar factor to match the absolute magnitude of the radio signal. The curve N/ρ (dotted) represents a properly normalised volume density of particles and appears to match the Charge Excess longitudinal evolution while N/ρ^2 (dashed) matching the Geomagnetic component could be interpreted as an additional factor of $1/\rho$ applied to the particle density.

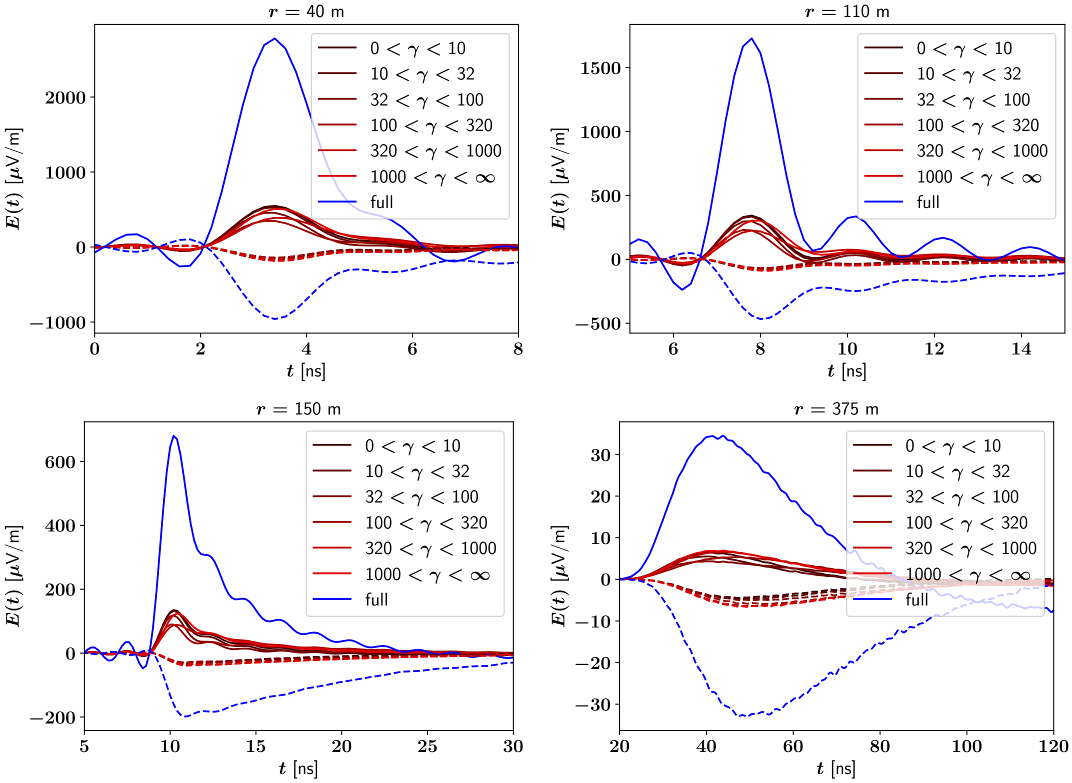

Figure 31: Signal at ground level sliced by the Lorentz factor γ of the emitting particles. The physically observable signal is shown in blue while the approximately logarithmic bins in Lorentz factor are shown in red and fade to black as γ decreases. For all lateral distances and both Geomagnetic (solid lines) and Charge Excess (dashed lines) emission all γ bins appear to contribute comparable amounts of radio emission.

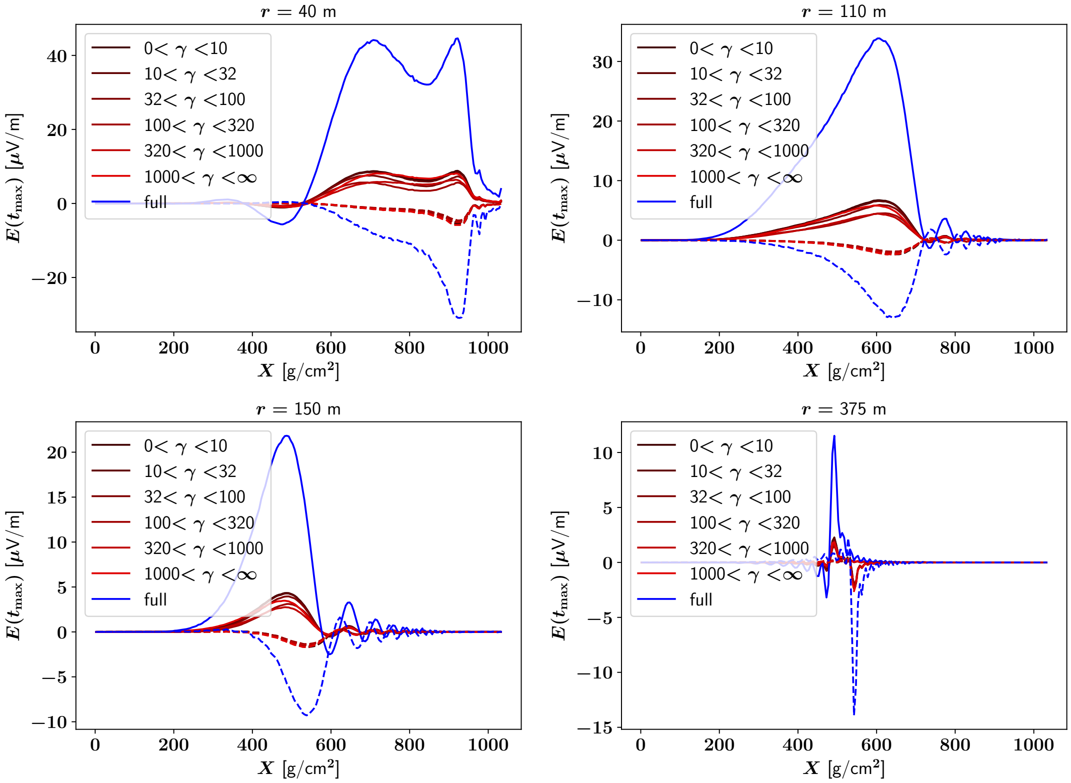

Figure 32: Value of the individual slice time series at the time t_{max} of the respective peak in the physical signal again split by the Lorentz factor γ of the emitting particles. The physically observable signal is shown in blue while the approximately logarithmic bins in Lorentz factor are shown in red and fade to black as γ decreases. Again for all lateral distances and both Geomagnetic (solid lines) and Charge Excess (dashed lines) emission all γ bins appear to contribute comparable amounts of radio emission to the observable peak.

interesting to investigate electron and positron populations of different energy individually in the future.

7

UNIVERSALITY CORRECTION FACTOR

Despite technical difficulties those configurations where the relative evolution-based synthesis did function showed clear improvements over our simplest model. This suggests our underlying assumptions may still be correct aside from the propagation effects discussed in section 6.5. In particular we clearly require an additional rescaling component to account for the evolution of particle distributions beyond their mere number. As discussed in section 5.4 the deviation of our synthesis result from the real signal appears to be directly proportional to the difference in X_{max} between the template and target showers. Instead of reconfiguring the entire template as in the previous section we now consider ways of directly compensating for this proportionality. Unfortunately this approach is entirely phenomenological and any physical origin will need to be determined retroactively once a functioning configuration is found.

7.1 CONCEPT

The first question to answer is whether our observation of differences between template and target being proportional to the difference in X_{max} as per figure 19 still holds on the level of individual slices. We can check this already using our slice configuration from sections 6.5 and 6.6 by normalising the radio signal to a reference point as shown in figure 33. Per construction the vertical separation of different slices shows the necessity of a longitudinal profile for rescaling while the low shower-to-shower fluctuations prove that adapting to the cascade evolution can work if done accurately enough.

The simplest approach is a parametric correction based entirely on the depth of the shower maximum X_{max} designed to account for the cascade evolution when added to our original model as per equation 9. When translated to our slicing at fixed atmospheric depth we expect this correction factor to have both longitudinal and lateral variations

$$\vec{E}^{\text{synth}}(r,t) = \sum_X \frac{U(X_{\text{max}}^{\text{real}}, r, X)}{U(X_{\text{max}}^{\text{temp}}, r, X)} \cdot \frac{N^{\text{real}}(X)}{N^{\text{temp}}(X)} \vec{E}^{\text{temp}}(\vec{r}, t, X) \tag{16}$$

where the parametric form of $U(X_{\text{max}}, r, X)$ must be continuous in X_{max} but may in principle be tabulated for our discrete set of slices in X and lateral positions r. Further an ideal construction would naturally converge to unity at large lateral distances and thus not alter our already near-perfect synthesis in this regime.

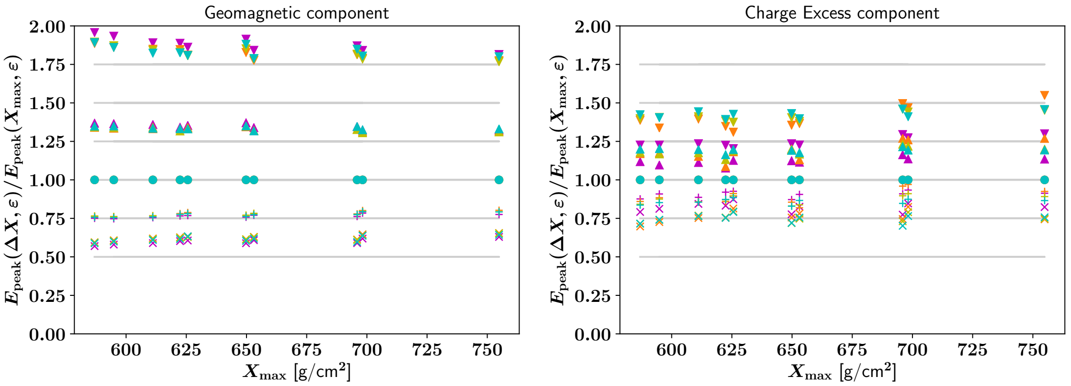

Figure 33: Estimate of longitudinal and shower-to-shower variations between different slices. Plotted are the ratios of peak amplitudes of the Geomagnetic (left) and Charge Excess (right) component for a slice at fixed ΔX relative to the respective value at X_{max}. Triangular markers denote slices above X_{max}, crosses those below and circles the reference point directly at X_{max} itself. The same values of ΔX are shown for 10 sample showers with varying X_{max} and four viewing angles of 0.5° (orange), 1° (yellow), 2° (cyan) and the Cherenkov angle (magenta). The grey lines serve only as a visual aid and all ratios for the slice at X_{max} are exactly equal to 1 by definition.

7.2 CONSTRUCTION VIA EMPIRICAL FIT

Ideally we would like a continuous analytic description but our templates are intrinsically limited to a finite number of sample points thus we can afford to leave this for proper theoretical work at a later date. In order to evaluate the feasibility of this approach and to obtain the correction model $U(X_{\mathrm{max}})$ we simulate a large number of air showers and plot the peak amplitude of the radio signal normalised with the particle number $N(X)$ for a given slice and antenna position over the depth of shower maximum for each sample as illustrated in figures 34 and 35. As expected some fluctuations remain but all physically relevant slices and positions show a clear correlation between the rescaled radio signal and X_{max}. The specific choice of our radio metric and the cascade evolution parameter on the axes does not qualitatively alter the result. For example the relative evolution stage or shower age may be used on the horizontal axis and the total energy in the pulse, the distance between the global minimum and maximum or the peak height of the Hilbert envelope on the vertical axis. For visual simplicity and numerical stability we opt to show a linear fit which may not be sufficient for all slices in particular when attempting to generalise to different primary particle species and energies. An overview and samples of the fit results are shown in figures 36 and 37,

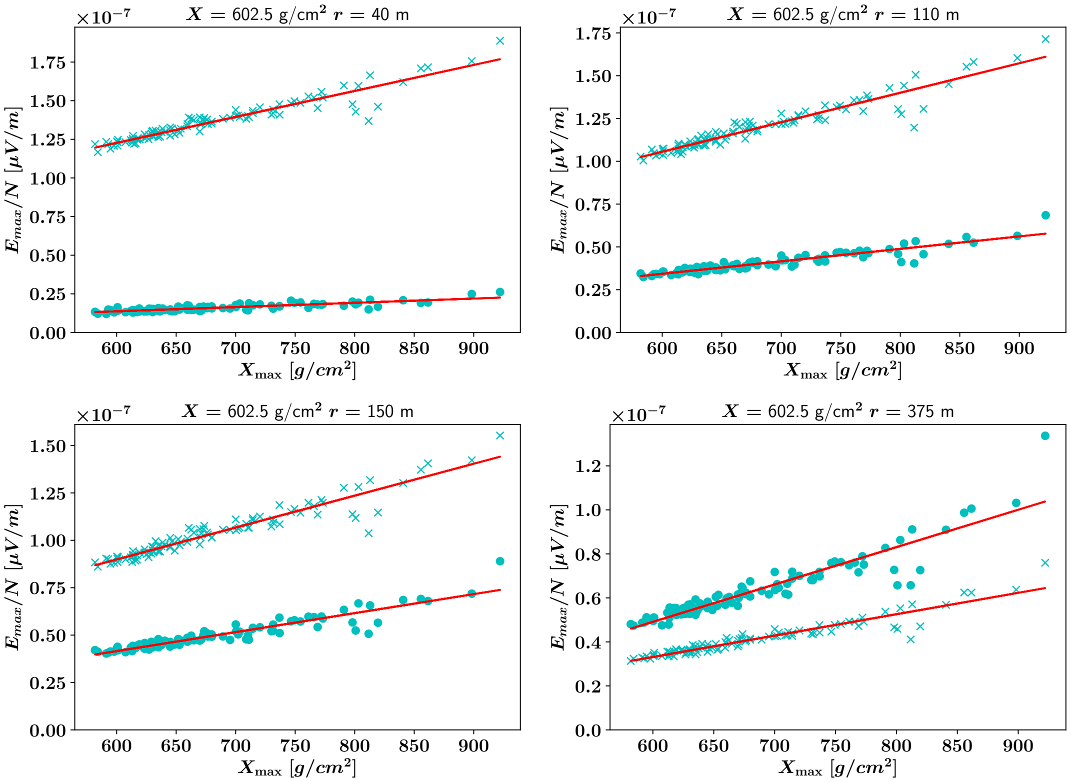

Figure 34: Correlation plots of the peak amplitude of the radio signal divided by particle number E_{max}/N plotted over the depth of the shower maximum X_{max} for a total of 100 vertical 10^{17} eV proton showers. Crosses denote Geomagnetic emission and circles the Charge Excess component. Shown here are four lateral distances one within the Cherenkov ring (top left), one close to it (top right), one outside (bottom left) and one far outside (bottom right) for a single slice near the top of the X_{max} spread range. The fitted linear correlation is plotted in red.

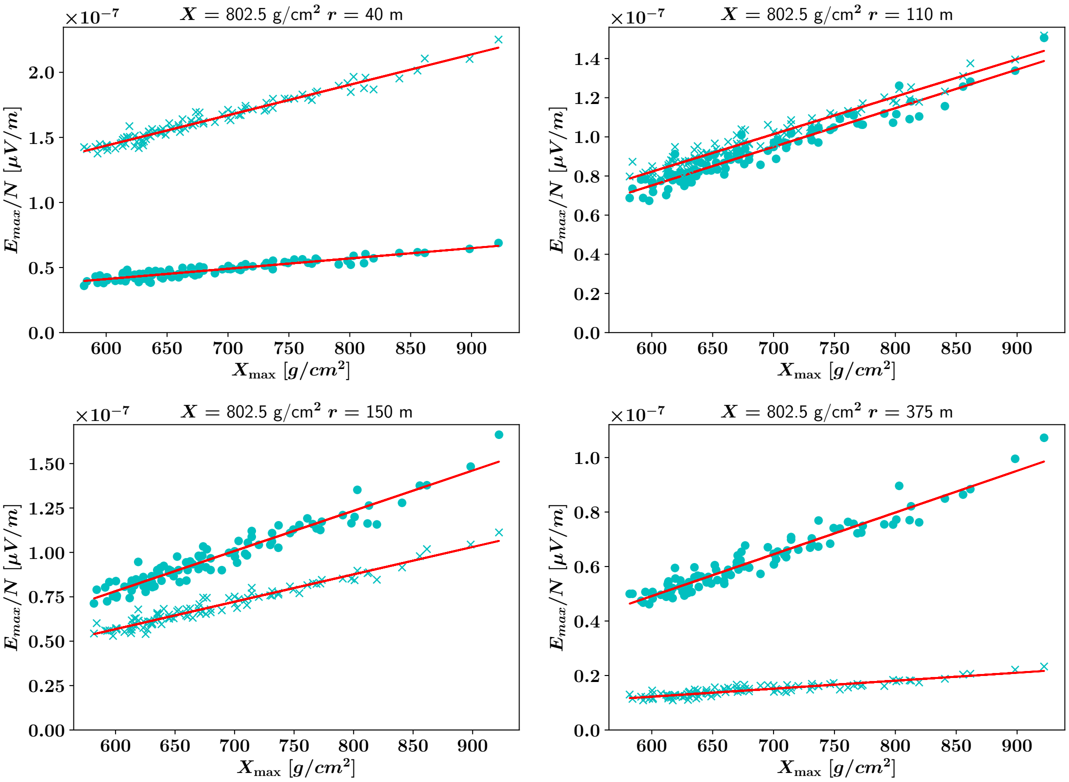

Figure 35: Same as figure 34 but for a slice near the end of the $X_{\max}$ range.

most importantly we do indeed see both longitudinal and lateral structure which does not match any simple curve shape.

7.3 APPLICATION AND LIMITATIONS

In testing this cascade-based correction factor $U(X_{\max})$ in synthesis some improvements immediately become apparent. The peak amplitudes within the Cherenkov ring can be matched almost perfectly even for large differences between the shower maxima. Several downsides are also highly visible however: at large lateral distances this approach performs considerably worse than either previous model despite matching the peak amplitude in individual contributing slices fairly well. More importantly the pulse shape is distorted with respect to the real signal at all lateral distances. All these observations are exemplified in figure 38. The mismatch in pulse shape proves a scalar correction **cannot** fully compensate the effects of the cascade evolution on the radio signal and we will have to investigate more complex solutions capable of distorting pulses in the following section.

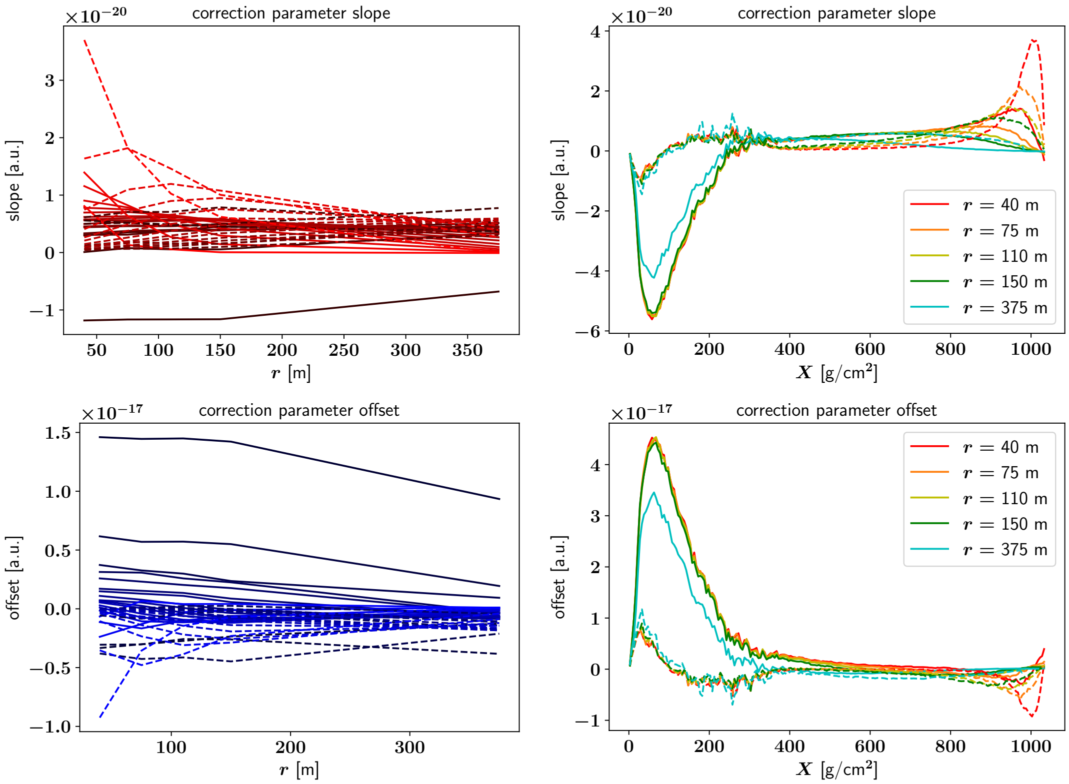

Figure 36: Lateral (left) and longitudinal (right) variation of the extracted correlation slope (top) and offset (bottom) parameters used in the fits exemplified by figures 34 and 35. Dashed lines indicate Charge Excess emission while the Geomagnetic component is plotted with solid lines. In the left-hand lateral plots multiple slices are plotted with colours fading to black with increasing atmospheric depth. In the left-hand longitudinal plots multiple lateral distances are shown in different colours. The very early and very late stages of longitudinal development are not reliable due to very low particle counts and increasing incoherence respectively but as shown in section 5.6 the range from ≈ 300 g/cm^2 to 950 g/cm^2 may contribute to observations.

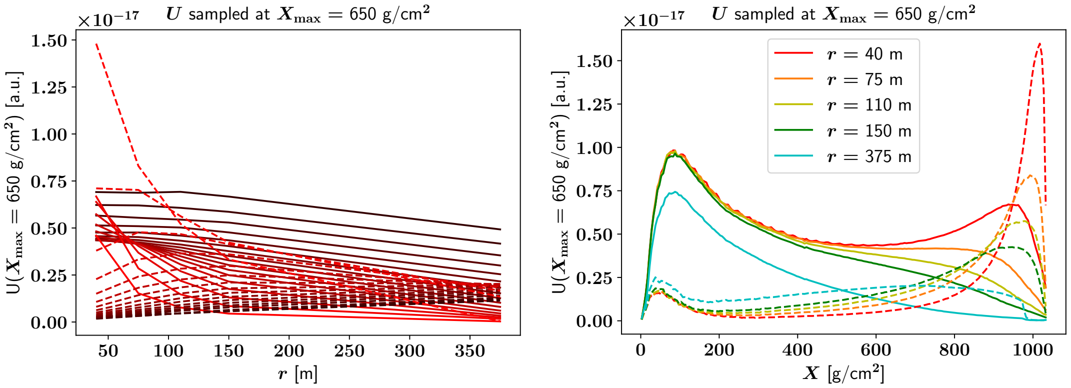

Figure 37: Sample of the correction factor $U(X_{\text{max}})$ itself for a "typical" X_{max} value of 650 g/cm^2. Again the lateral variation for multiple slices is shown to the left and the longitudinal evolution for several lateral distances to the right using the same lines and colours as in figure 36.

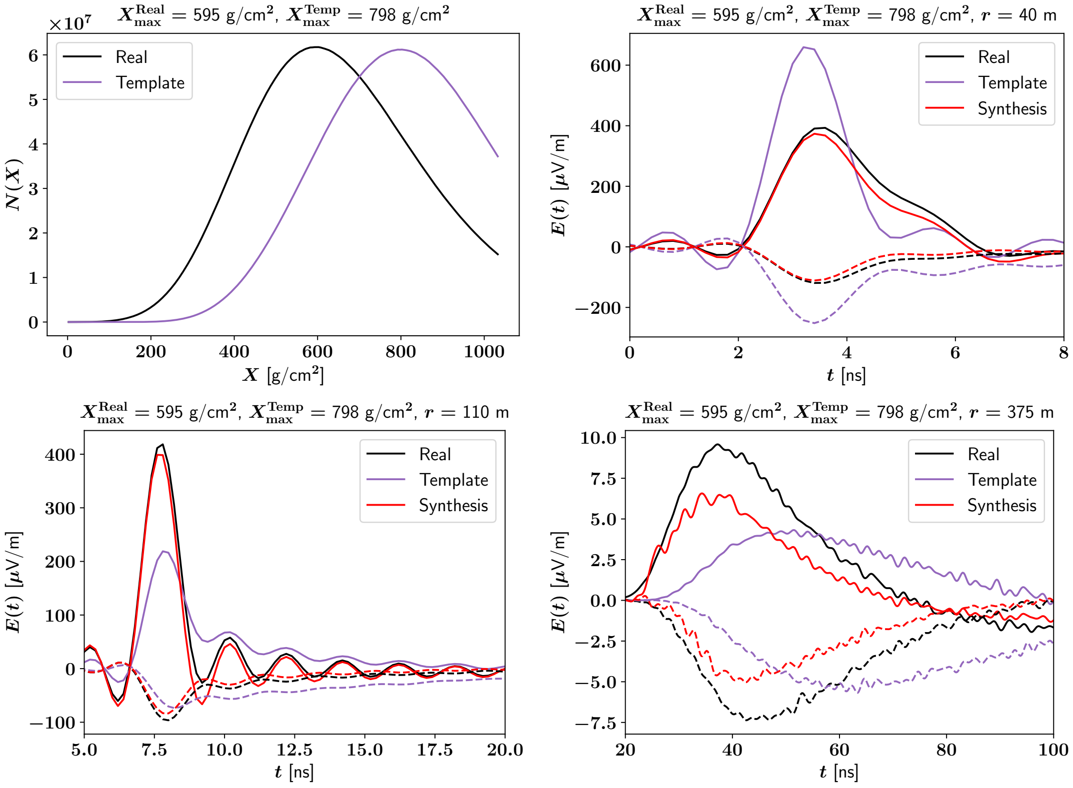

Figure 38: Longitudinal profiles (top left) and synthesised time series within (top right), near (bottom left) and outside (bottom right) the Cherenkov ring using a scalar X_{max}-dependent correction factor determined from empirical correlations. While the peak amplitudes at 40 m and 110 m show clear improvements over previous synthesis models the pulse shape (e.g. the first minimum after the peak at 110 m) is reproduced *worse* than by previous models. At 375 m the correction fails entirely matching neither the pulse shape nor the peak amplitude it was designed to normalise for individual slices.

8

ACCURATE REFINEMENT - SPECTRAL RESCALING

A scalar factor fundamentally cannot alter the pulse shape. Therefore we need to apply a higher order rescaling which relates the depth of the shower maximum to features in the pulse structure of each individual slice time series.

8.1 AMPLITUDE FITTING

In order to extend our correction model further we choose to consider the full frequency spectrum of each individual slice time series as the very narrow width of the pulses in the time domain is prone to amplifying numerical errors. The phase spectra ought to be dominated by a linear gradient representing the arrival time of the pulse as per our arrival time model thus we focus on the amplitude spectra. Such spectral descriptions are common for physically observable pulses but to our knowledge a precise investigation of individual contributing slices is entirely unprecedented in the context of cosmic ray air showers therefore we are limited to an empirical treatment based on the dominant visual features of the spectra. Following the longitudinal evolution of a single shower for a fixed antenna position we see a clear and mostly continuous evolution of the amplitude spectrum from each individual slice to the next. This we interpret as a compound effect resulting from both the changing viewing angle and propagation geometry as well as an evolution of the emitting particle distribution. The spectra may display some or all of the following features:

- **Low-frequency cutoff:** prominent for earlier stages and late-developing showers the spectrum drops rapidly towards the lowest frequencies. Most of this range is not accessible in measurements due to omnipresent atmospheric and galactic backgrounds [49] thus a more gentle curve may be sufficient for practical applications.

- **Smooth curve:** the bulk of the coherent radio signal's power. Depending on other features it may be linear or display a slight curve adequately described by a quadratic exponential.

- **High-frequency attenuation:** primarily relevant for the very late stages of the shower but also adding to overall curvature the coherent portion of the radio signal diminishes with increasing frequency. The location of the visual cutoff point varies from shower to shower based on X_{max} while its shape appears to be dominated by geometry.

- **Incoherent noise:** a combination of physically real incoherent signal contributions and numerical artifacts within the framework of CORSIKA+CoREAS. Fluctuations around a non-zero mean pose no challenge but at high frequencies the spectrum contains *only* noise which may compromise spectral fits.

Accurately treating the low-frequency cutoff leads to substantial numerical instability making it non-viable for largely unsupervised automation in the following steps. As simultaneous occurrences of the low-frequency cutoff and full high-frequency attenuation are incredibly rare we opted to fit a single quadratic function to each spectrum individually with the role of the quadratic term shifting accordingly. An additional term became necessary to buffer the noise floor which would otherwise destabilise the fit. Not wanting to risk degeneracies with very flat slopes the noise floor d was loosely parametrised in advance leaving three spectral coefficients A_0, b and c

$$\tilde{A}_{\text{slice}}(f) = A_0 \cdot \exp(b \cdot f + c \cdot f^2) + d \tag{17}$$

with

$$\sqrt{d} = \max\left[10^{-9} \cdot \left(\frac{X}{400\text{g/cm}^2} - 1.5\right) \cdot \exp\left(1 - \frac{r}{40000\text{cm}}\right), 0\right] \tag{18}$$

tuned by eye to roughly match the total incoherent noise floor across all lateral distances and primarily for the later stages of shower development where the coherent spectrum cuts off within our frequency band. Example applications of this fit are shown in figure 39. Each slice and antenna position are still fitted independently from one another meaning this approach can fully recover the previous scalar correction by setting $b = c = 0$ if necessary.

8.2 AMPLITUDE SPECTRUM CORRECTION

Having reduced the full amplitude spectra to three scalar parameters we can once again draw correlation plots for many sample showers. Currently we fit the X_{max} correlations of the three spectral parameters separately and also continue to do so independently for each slice and antenna position. As shown in figure 40 a linear fit suffices once more although the exponential coefficients scatter significantly. By eye we can also tell a single slightly curved fit can capture all primary cosmic ray species and energies. This observation is crucial for a long-term generalisation and application of our model as it allows us to operate from a single set of spectral correction parameters which require many samples to determine. Hailing back to section 4.2 these correlations also imply a high degree of universality in the radio signal tied entirely to the longitudinal profile as X_{max} is the only shower-dependent input to the

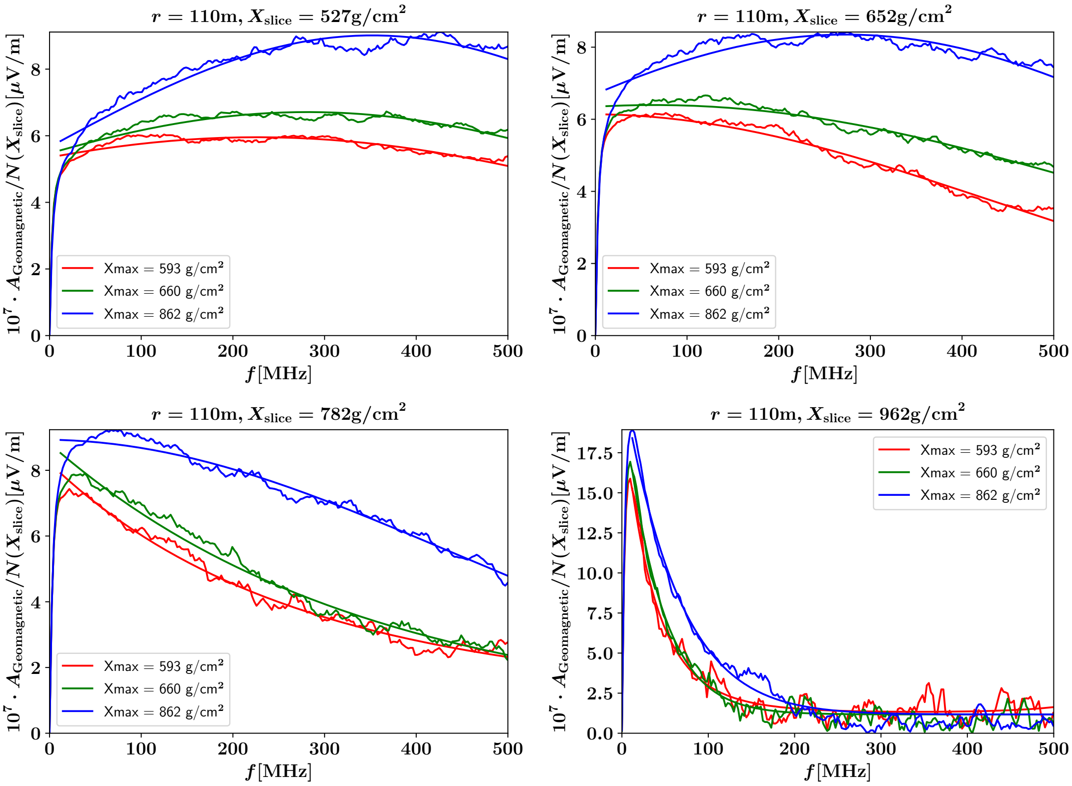

Figure 39: Fits of amplitude spectra for slices near the beginning (top left), centre (top right) and end (bottom left) of the typical range of X_{max} values for the vertical 10^{17} eV protons initiating these showers as well as an example of the noise- or incoherence-dominated very late evolution (bottom right). Three sample showers with very early (red), "typical" (green) and late (blue) shower maxima are shown in each plot. All spectra are from the same antenna located near the Cherenkov ring. The hierarchical ordering by X_{max} is consistent for other lateral distances as well.

fit. Examples using a quadratic function are shown in figure 41 however with the bulk of the shower maxima being cocentrated around "typical" values we avoid using such curved fits when explicitly investigating differences in $X_{\max}$ where the larger values represent rare edge cases with little weight in the correlation fit. The parameters b and c which enter the exponential in equation 17 in particular may cause explosive instabilities or even numerical overflows when comparing the most extreme ends of the correlation curve (e.g. early-developing 10^{17} eV iron showers versus deeply penetrating 10^{19} eV protons).

The individual coefficients of the polynomial fits now form the additional parametrised modification of our synthesis procedure. Each parameter of the amplitude spectrum can be predicted for a target shower given only the location of its longitudinal maximum $X_{\max}$. Computation speed and the memory demands of the coefficient table are sufficiently low as well. In figure 42 we show the longitudinal variation of the resulting spectral parameters. The individual fit coefficients show similar curves which we skip for simplicity. We have not been able to determine which features in this longitudinal variation have a physical origin but can tell with relative certainty from inspecting individual amplitude spectra that the very early stages at $X < 200$ g/cm^2 and the very late stages at $X > 950$g/cm^2 are completely incoherent. These regions do not contribute significantly to the radio signal however therefore we merely modify the resulting amplitude spectra in template construction and synthesis to avoid purely numerical errors such as division by near zero.

Fully unraveling the physical origins of the amplitude spectrum scaling would require significant theory work as well as further dedicated simulations. Due to time constraints we instead focus on completing a functional core synthesis model.

8.3 FINAL SYNTHESIS MODEL

Here we give an overview of the full synthesis procedure for a single antenna coinciding with a template location. While the individual steps are identical we now separate the template processing and the actual synthesis. Our criterion for this distinction being that every aspect of the template processing can be performed once without knowing any future synthesis inputs and the results stored without significantly increasing the file size compared to raw Monte Carlo outputs and parametrisation tables. Every step classified as part of the synthesis meanwhile requires knowledge of at least some physically relevant quantities of the individual target shower.

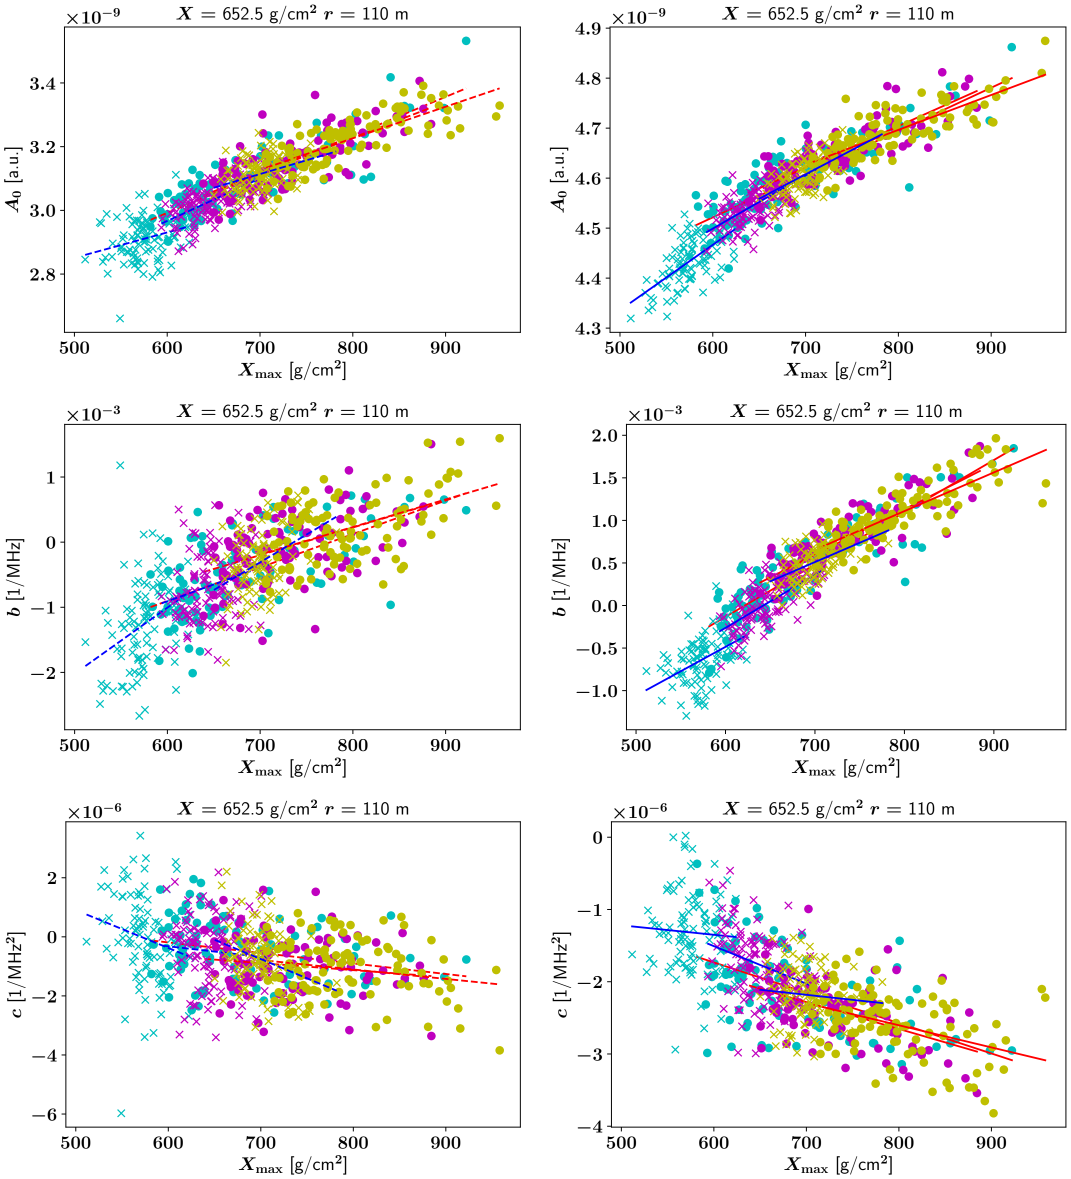

Figure 40: Correlation fits for the slice at 650 g/cm^2 which is a “typical” value for $X_{\max}$ and an antenna near the Cherenkov ring. Each plot shows showers of both iron (crosses) and proton (circles) primaries with three fixed energies of 10^{17} eV (cyan), 10^{18} eV (magenta) and 10^{19} eV (yellow) each represented by 100 showers of varying $X_{\max}$. The plots to the left show Charge Excess emission, those to the right the Geomagnetic component. All three coefficients A_0 (top), b (middle) and c (bottom) are fitted completely independently from one another, in this case the two primary species and three primary energies were each also fitted independently with a straight line (blue for iron, red for protons).

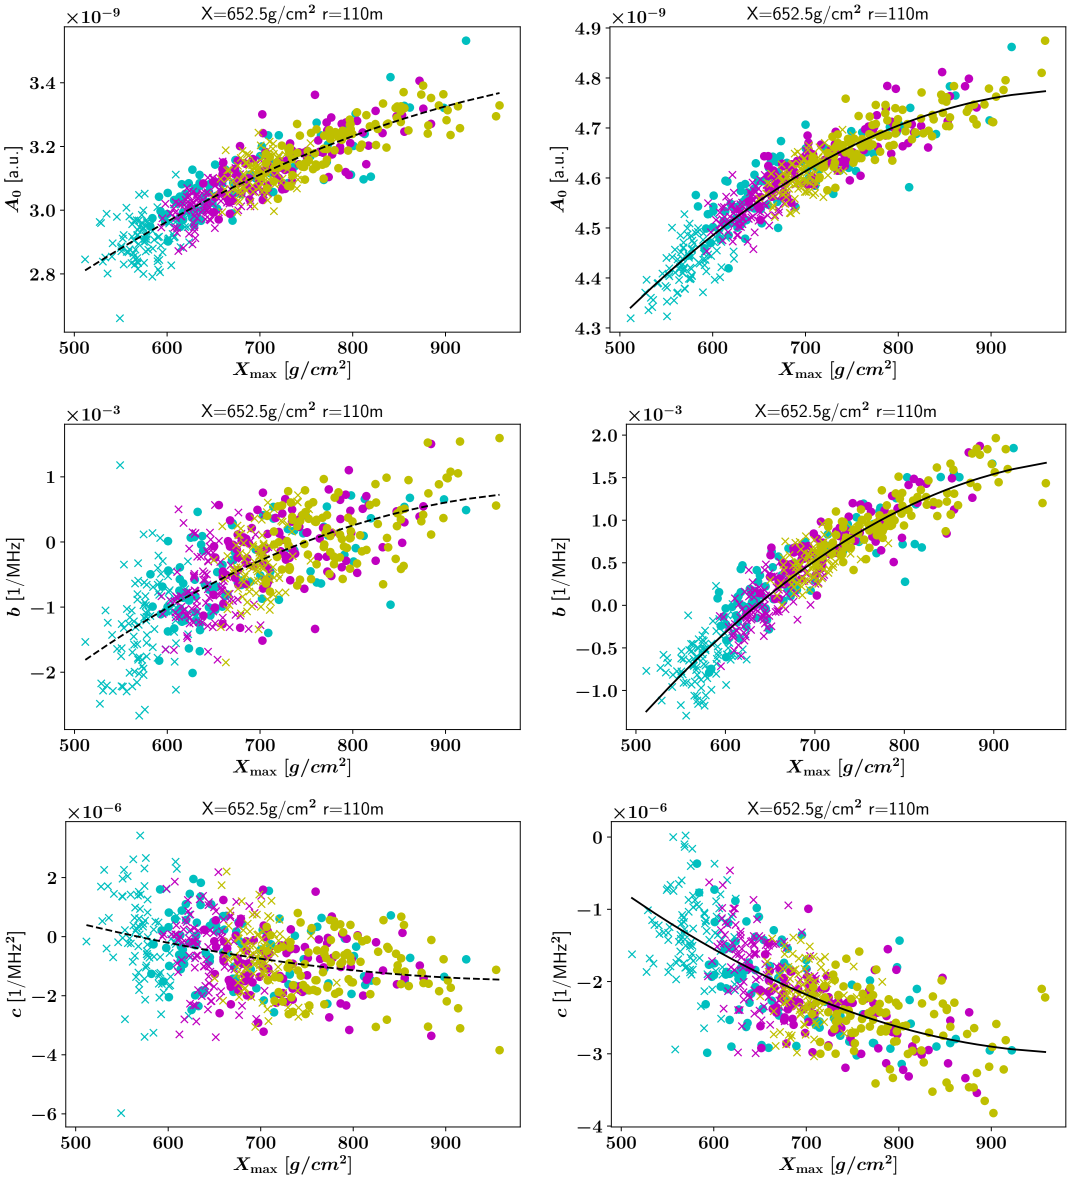

Figure 41: Same as figure 40 but now fitting all primary energies (colours) and species (markers) combined with a quadratic function (black).

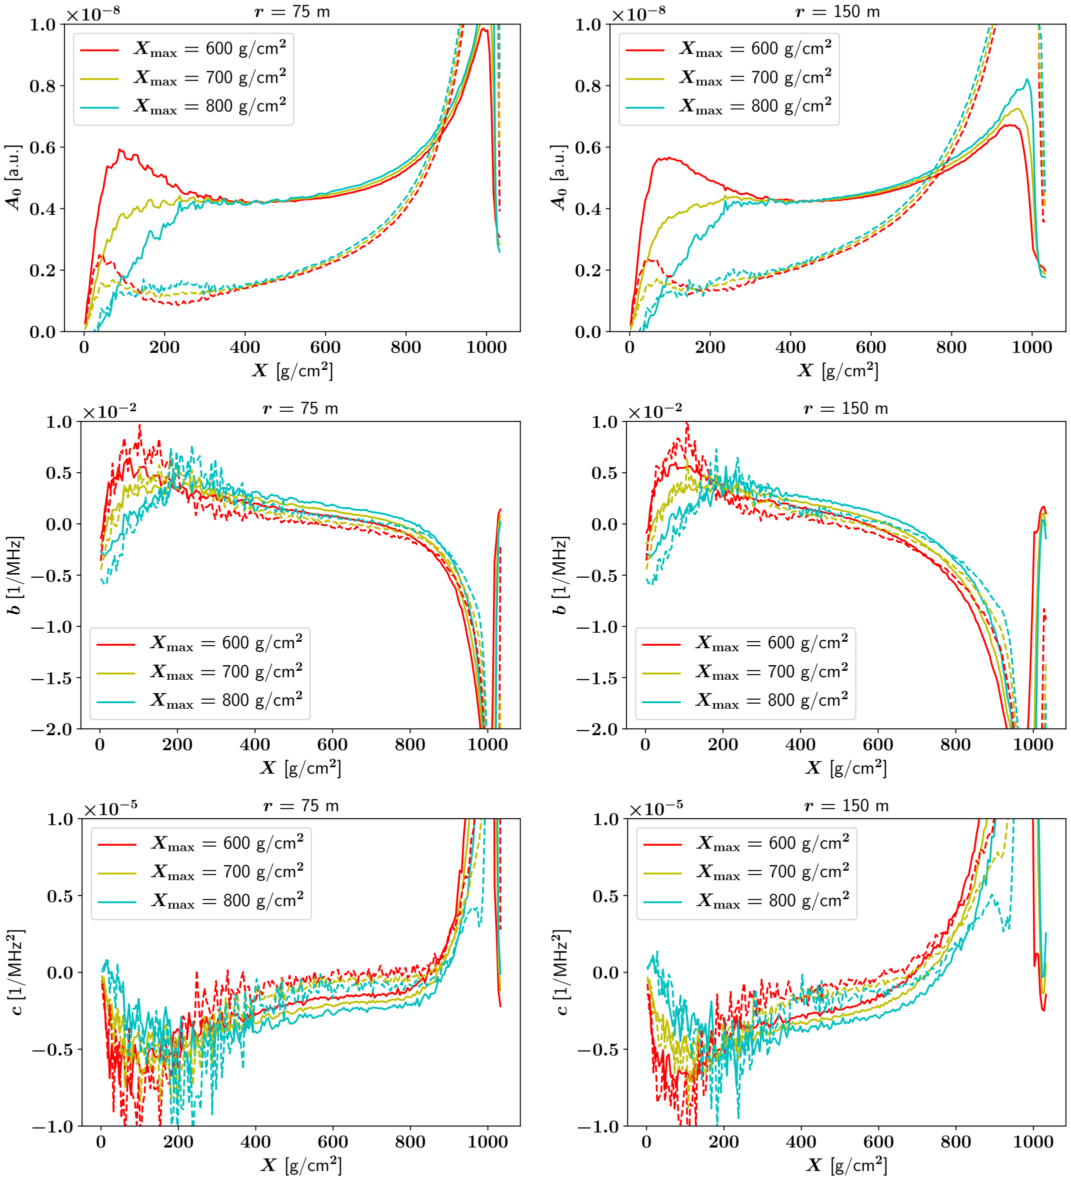

Figure 42: Longitudinal evolution of the spectral parameters A_0 (top), b (middle) and c (bottom) for two lateral distances inside (left) and outside (right) the Cherenkov ring. Three values of $X_{\max}$ were given to evaluate the polynomial correlation fit. The Geomagnetic component (solid lines) and Charge Excess emission (dashed lines) are treated independently starting from the original amplitude spectrum fit.

8.3.1 *Template Processing*

In order to create a template for our synthesis model we first require a full CoREAS simulation of a single line of antennas along the $\vec{v} \times \vec{v} \times \vec{B}$ axis and each of these positions sliced longitudinally using the existing CoREAS options. Ideally one would compute the template shower without any numerical thinning, if not feasible it is still recommended to thin less than for normal "production" runs: CORSIKA+CoREAS simulations typically engage thinning at a fraction of 10^{-6} relative to the primary energy while we observed differences when reducing the threshold to 10^{-7}. Note computation time does not scale linearly with the thinning threshold: conservation of energy allows far more low-energy particles to be produced from the same budget. Despite our goal of a universal template it may be desirable to select several with vastly differing shower maxima as almost all observed fluctuations scale with the difference between template and target depth of shower maximum. If only one template is used it should be pre-selected to have a fairly (but not statistically exceptional) early X_{max} in order to encompass as much of the late stages of cascade development as possible and to reduce the impact of near-field effects. Determination of the amplitude spectrum correction parameters requires a larger study of hundreds of showers which should represent the target primary particles and also match the thinning level of the template simulations.

We still begin with the longitudinal profile as in equation 9 which may also be substituted with a reliable Gaisser-Hillas fit as per equation 4

$$\Gamma_{\text{basic}}^{\text{temp}} = N^{\text{temp}}(X) \tag{19}$$

while Fourier-transforming all raw CoREAS time series

$$\vec{E}^{\text{temp}}(\vec{r}, f, X) = \mathcal{F}\left[\vec{E}^{\text{temp}}(\vec{r}, t, X)\right] \tag{20}$$

and expressing each polarisation individually in polar form

$$E^{\text{temp}}(\vec{r}, f, X) = A(\vec{r}, f, X) \cdot \exp\left(i \cdot \phi(\vec{r}, f, X)\right) \tag{21}$$

while computing the amplitude spectrum correction. In principle one could fit each amplitude spectrum directly but this is more volatile and requires greater care compared to using the parametrisation averaged over many different showers therefore we evaluate the fitted polynomials

$$A_0^{\text{temp}} = p_0^A + p_1^A \cdot X_{\text{max}}^{\text{temp}} + p_2^A \cdot \left(X_{\text{max}}^{\text{temp}}\right)^2 \tag{22}$$

$$b^{\text{temp}} = p_0^b + p_1^b \cdot X_{\text{max}}^{\text{temp}} + p_2^b \cdot \left(X_{\text{max}}^{\text{temp}}\right)^2 \tag{23}$$

$$c^{\text{temp}} = p_0^c + p_1^c \cdot X_{\text{max}}^{\text{temp}} + p_2^c \cdot \left(X_{\text{max}}^{\text{temp}}\right)^2 \tag{24}$$

where the quadratic coefficients p_2^{par} are zero when fitting each primary species and energy separately. From these spectral parameters we calculate the amplitude spectrum model as per equation 17

$$\tilde{A}^{\text{temp}}(\vec{r}, f, X) = A_0^{\text{temp}} \cdot \exp(b^{\text{temp}} \cdot f + c^{\text{temp}} \cdot f^2) + d \tag{25}$$

$$\Gamma_{\text{amp}}^{\text{temp}}(\vec{r}, f, X, X_{\text{max}}^{\text{temp}}) = \left[\tilde{A}^{\text{temp}}(\vec{r}, f, X)\right]^{-1} \tag{26}$$

where if d is not calculated in accordance with equation 18 or a similar model the spectrum should still be clamped to a reasonable constant value as the fitted exponential can drop too rapidly and even approach numerical zero within the frequency range of interest. The inverse of the amplitude spectrum fit forms our spectral correction factor to be applied individually for each slice, antenna position and polarisation. If any kind of arrival time manipulation is desired there are several options. As the normalised template will initially exist in the frequency domain this also presents an ideal opportunity to perform any timing shifts in the form of adding a linear gradient to the phase spectrum $\phi(f)$. Despite being primarily designed for time-domain output the nature of our amplitude spectrum fit makes it more efficient to store the processed templates in the frequency domain. If future improvements allow complete omission of the phase spectra this could potentially halve the space requirements for template storage but we currently save both them and the amplitude residuals. Thus each stored template consists of

$$A_{\text{res}}(\vec{r}, f, X) = \Gamma_{\text{basic}}^{\text{temp}}(X) \cdot \Gamma_{\text{amp}}^{\text{temp}}(\vec{r}, f, X, X_{\text{max}}^{\text{temp}}) \cdot A(\vec{r}, f, X) \tag{27}$$

$$\phi_{\text{res}}(\vec{r}, f, X) = \phi(\vec{r}, f, X). \tag{28}$$

For ease of access we also store the global time reference of each time series explicitly although in principle the physically meaningful time differences should be encoded in the phase spectra. In principle explicit knowledge of the template shower's longitudinal profile should not be required from this point onwards but if using multiple templates the depth of shower maximum X_{max} should be the primary identifier.

8.3.2 *Synthesis Runtime*

For small-scale simulations the bulk of our runtime may be spent on initially loading the templates which remains significantly more demanding than analytical models. Thus grouping larger sets of simulations based on the properties of the primary particle and especially its arrival direction may be necessary to reduce loading times and memory demand. For each individual shower we require only the physical properties of the primary cosmic ray and the longitudinal profile of electrons and positrons $N^{\text{real}}(X)$. In most cases this longitudinal profile is accurately represented by a Gaisser-Hillas fit which can be used to interpolate mismatched longitudinal resolutions as well as provide a smooth analytical prediction of X_{max}. The longitudinal profile sampled at the template slice resolution forms our basic scaling factor

$$\Gamma^{\text{synth}}_{\text{basic}} = N^{\text{real}}(X) \tag{29}$$

while the amplitude spectrum coefficients are evaluated with the same polynomial fits as for the template

$$A_0^{\text{synth}} = p_0^A + p_1^A \cdot X_{\text{max}}^{\text{real}} + p_2^A \cdot \left(X_{\text{max}}^{\text{real}}\right)^2 \tag{30}$$

$$b^{\text{synth}} = p_0^b + p_1^b \cdot X_{\text{max}}^{\text{real}} + p_2^b \cdot \left(X_{\text{max}}^{\text{real}}\right)^2 \tag{31}$$

$$c^{\text{synth}} = p_0^c + p_1^c \cdot X_{\text{max}}^{\text{real}} + p_2^c \cdot \left(X_{\text{max}}^{\text{real}}\right)^2 . \tag{32}$$

Contrary to the template processing it is both physically realistic and desirable to omit the noise term d as our model does not treat incoherent signal contributions. This also prevents the amplitude correction terms from canceling to numerical unity when both the template and target spectrum fits drop to numerical zero within the frequency range of interest

$$\Gamma^{\text{synth}}_{\text{amp}}(\vec{r}, f, X, X_{\text{max}}^{\text{real}}) = A_0^{\text{synth}} \cdot \exp(b^{\text{synth}} \cdot f + c^{\text{synth}} \cdot f^2). \tag{33}$$

The two rescaling factors are then applied to the template amplitude data of a single slice in the frequency domain

$$A^{\text{synth}}(\vec{r}, f, X) = \Gamma^{\text{synth}}_{\text{basic}}(X) \cdot \Gamma^{\text{synth}}_{\text{amp}}(\vec{r}, f, X, X_{\text{max}}^{\text{real}}) \cdot A_{\text{res}}^{\text{temp}}(\vec{r}, f, X) \tag{34}$$

which is combined with the phase template

$$\vec{E}^{\text{synth}}(\vec{r}, f, X) = A^{\text{synth}}(\vec{r}, f, X) \cdot \exp\left(i \cdot \phi_{\text{res}}^{\text{temp}}(\vec{r}, f, X)\right). \tag{35}$$

At this point one could potentially insert additional processing steps such as shifting the phases or interpolating between antenna positions in close proximity to one another however these go beyond the scope of our current model. Finally the synthesised radio pulse is constructed in the time domain by summing up the contributions from all longitudinal slices after Fourier transforming each of them individually

$$\vec{E}^{\text{synth}}(\vec{r}, t) = \sum_X \mathcal{F}^{-1}\left[\vec{E}^{\text{synth}}(\vec{r}, f, X)\right]. \tag{36}$$

Due to the linearity of the Fourier transform our simple core model could perform the sum first and immediately return the synthesised radio signal in the frequency domain instead. We primarily convert intermediate results to the time domain for direct comparison with the original CoREAS time series as well as the general ease of monitoring progress. For future practical application the necessity of slice-by-slice Fourier transforms will depend on which generalisations prove most reliable in covering the full parameter space relevant to observations.

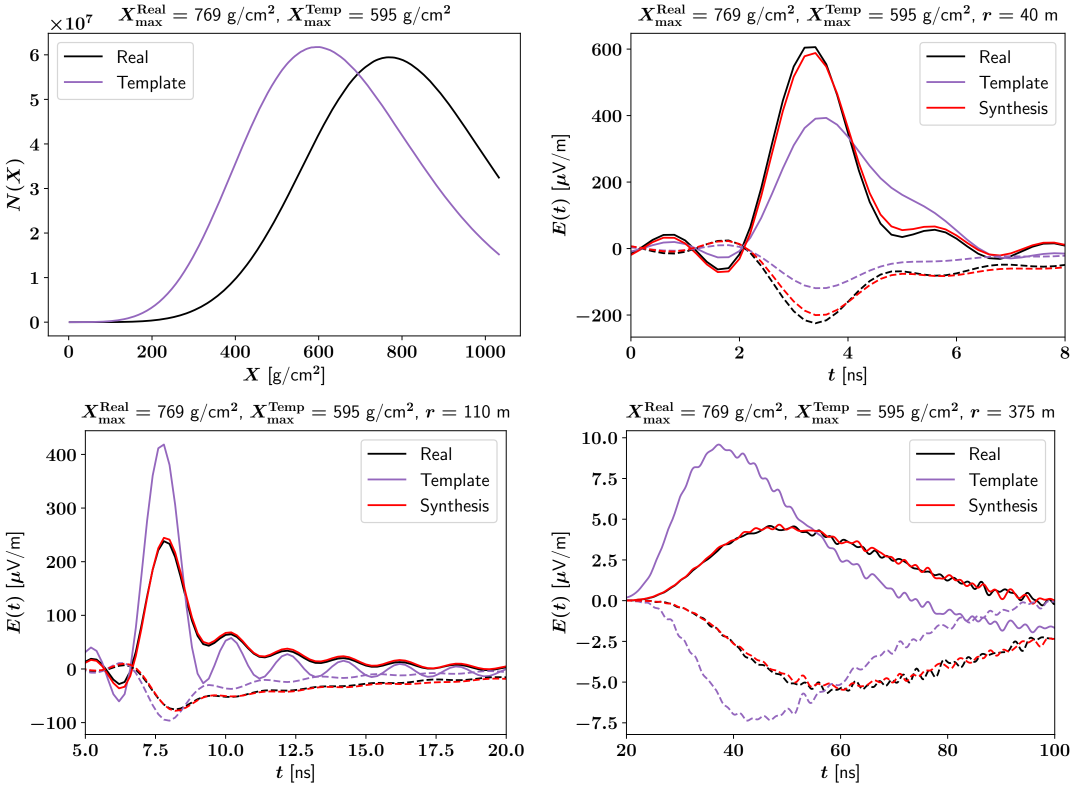

Figure 43: Amplitude-corrected synthesis for vertical 10^{17} eV protons for showers with a substantial $\Delta X_{\max}$ of -174 g/cm^2. As in previous figures we show the longitudinal profiles in the top left and then three lateral distances within (top right), near (bottom left) and outside (bottom right) the Cherenkov ring. Charge Excess emission is plotted with dashed lines and the Geomagnetic component with solid lines. Only slight deviations are visible and only for the smallest lateral distance.

8.3.3 *Results*

As shown in figure 43 this approach can achieve near-perfect synthesis accuracy even for shower maxima further apart than those which stumped our previous models. Some fluctuations are still expected as per our correlation plots in figure 40 so this example represents the ideal case where both showers do not deviate much from the correlation fit line. The combined fit across all primary species and energies sacrifices some accuracy as illustrated in figure 44 but remains quite impressive even for different choices of template and target primary cosmic rays. A detailed discussion of our results and the remaining limitations will be given in the following chapter 9.

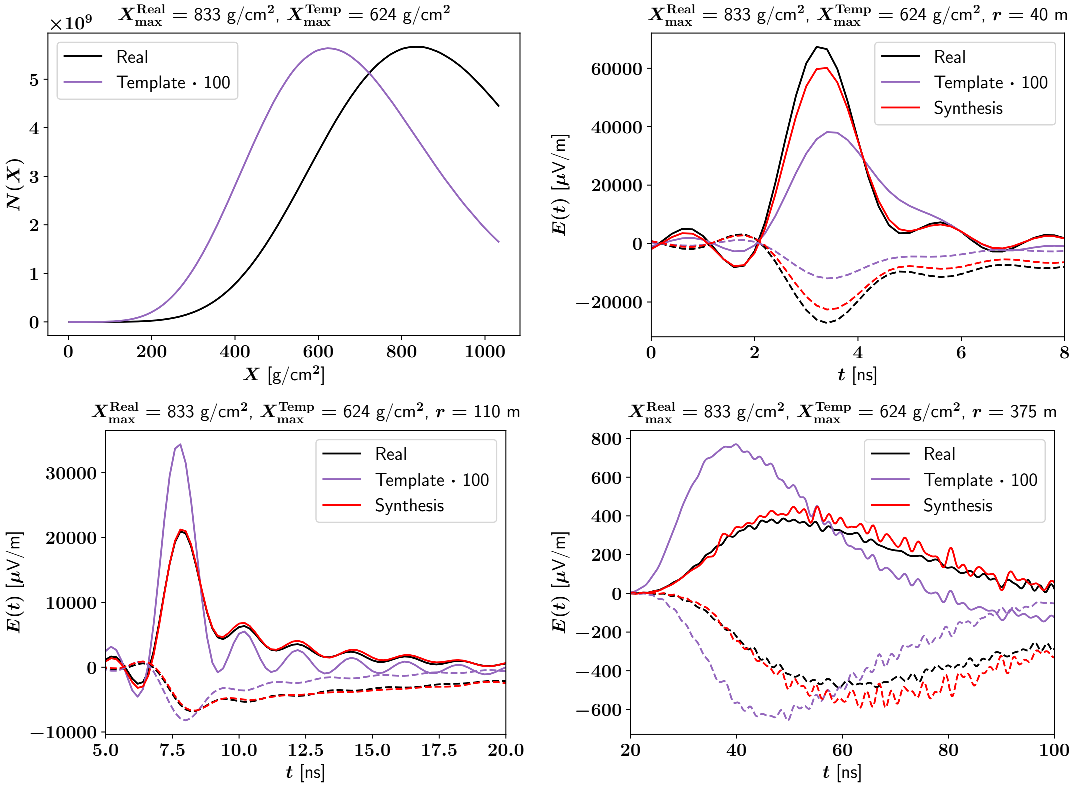

Figure 44: Amplitude-corrected synthesis for a 10^{17} eV iron-induced template and a 10^{19} eV proton target with a ΔX_{max} of -209 g/cm^2. The closest (top right) and furthest (bottom right) lateral distances show some deviations suggesting there may be room for further improvement.

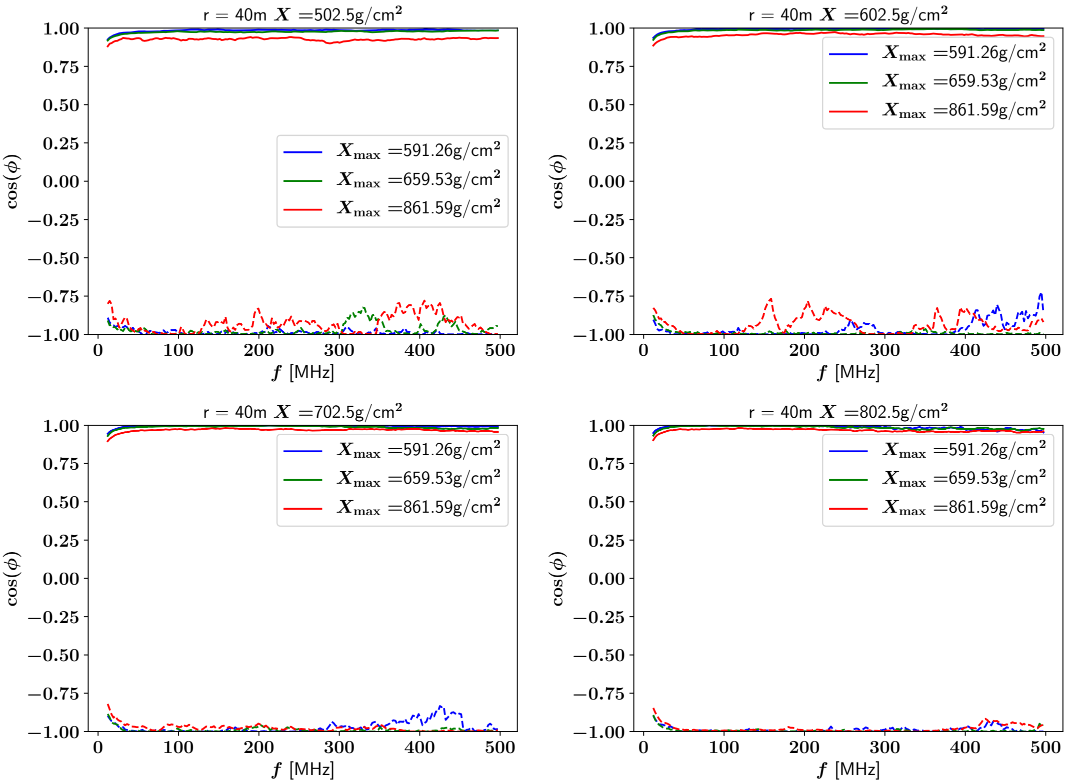

Figure 45: Cosine of the phase spectrum of individual slice time series, after correcting for the arrival time. The three colours correspond to three showers with equal primary particles (vertical 10^{17} eV protons) but vastly different longitudinal profiles. The Geomagnetic component is plotted with solid lines, the Charge Excess component with dashed lines. This lateral distance lies much closer to the shower core than the Cherenkov ring.

8.4 RESIDUAL PHASES

The observed success of our amplitude-based rescaling model implies a high degree of universality in the remaining phase spectra $\phi(f)$. Ideally they would be consistent with a linear gradient solely encoding the arrival time of the main peak which would allow us to omit the phase spectra from our templates entirely and only store those scalar arrival times instead. Indeed for small lateral distances as exemplified by figure 45 subtracting the physical arrival time of the pulse leaves a spectrum which is both highly universal and flat over the full range of relevant slices implying near-perfect accuracy of our timing model.

Closer to the Cherenkov ring e.g. in figure 46 we see a nonzero slope and minor fluctuations but the spectra of different showers remain very close to each other with the exception of the highest frequencies for slices very late in

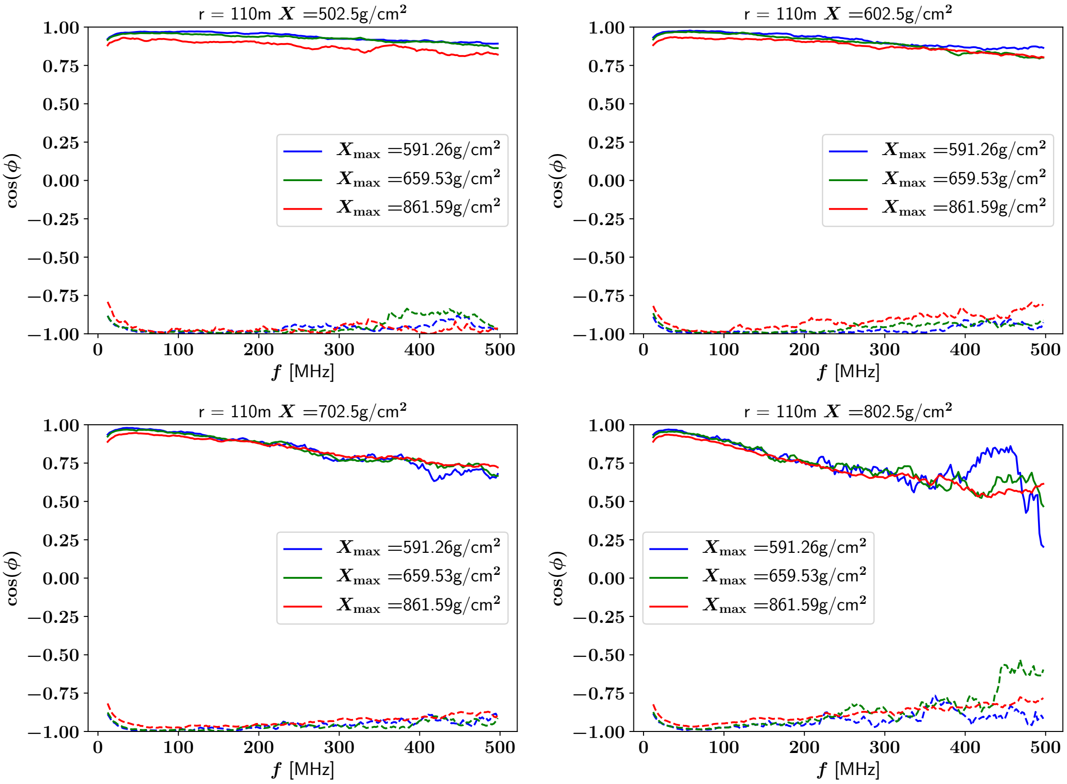

Figure 46: Same as figure 45 but for a lateral distance close to the Cherenkov ring. The phase spectrum still appears to be quite universal but contains additional structure not represented in our simple arrival time model. The very late development which shows shower-to-shower variations is unlikely to contribute due to the absolute amplitude of these slices.

the cascade evolution. Those slices are not expected to contribute but unless a fully reliable model of the average gradient is found we cannot abandon the use of phase templates entirely.

Finally for large lateral distances we see a wide variety of behaviours as shown in figure 47. Even for slices near the shower maxima a linear gradient no longer describes the residual phase spectra. Due to the time structure of the pulse more or less being a direct projection of the longitudinal cascade development with no concurrent arrivals or inverted timings we cannot easily tell which slices genuinely do not contribute to the observable pulse. Indeed that may largely depend on the observed bandwidth and detection thresholds as these broad pulses would slowly drop below the noise floor in real measurements. Due to fundamental limits of causality and propagation speed we can however be certain the very late stages of the cascade do not contribute any signal within realistic timing windows therefore the complete incoherence of the phase residuals does not impede our model.

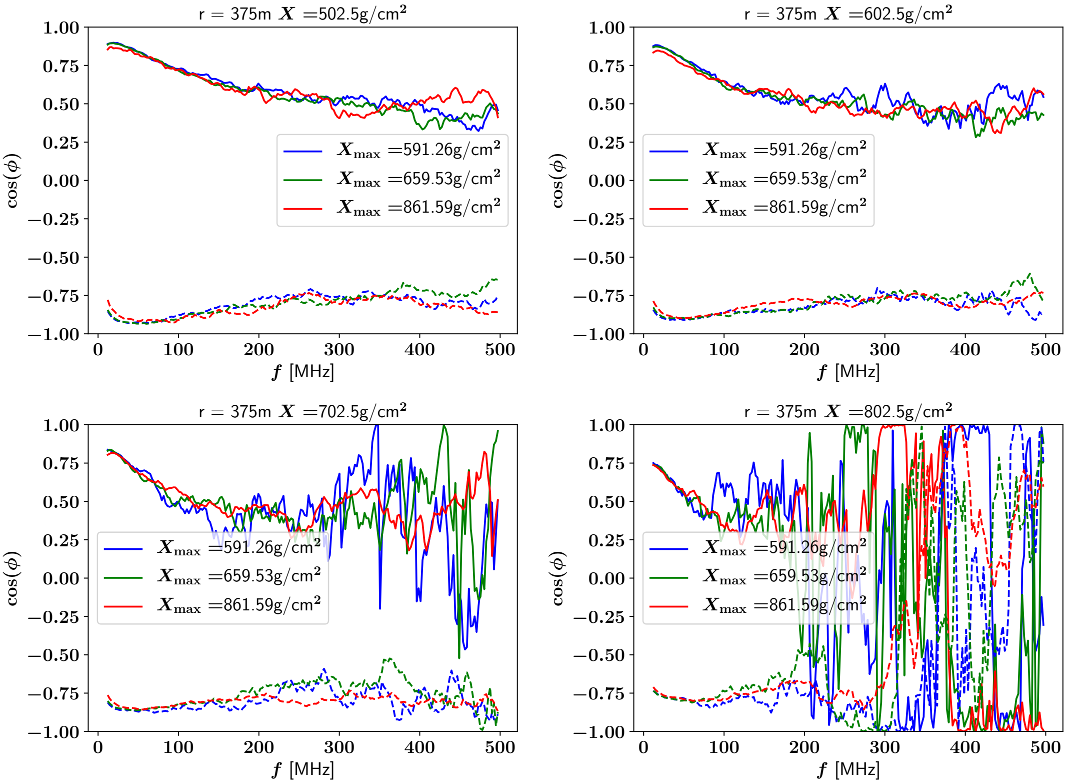

Figure 47: Same as figure 45 but for a lateral distance far beyond the Cherenkov ring. By comparison with figure 22 the incredibly unstable phases shown here contribute absolutely nothing to the commonly observed radio signal which we also confirmed on the time series level.

Despite the emergence of additional features we still observe a very good degree of shower-to-shower universality until the complete breakdown of coherence occurs. This suggests our model should remain accurate as long as the residual phases are preserved.

9

COMPARISON AND BENCHMARKING

In evaluating the performance of our synthesis procedure there are two primary deciding factors: accuracy and speed. Ideally we would like to match the quality of CoREAS simulations as the most accurate representation of the true physics we have access to. As these true cascade physics include very real statistical fluctuations caused by high-energy hadronic interactions beyond even the LHC centre-of-mass energy perfect equality can only be achieved by luck rather than correct parametrisation. Merely quantifying the statistical floor of higher order fluctuations present in CoREAS even when the longitudinal profile is fixed as described in section 4.2 across the full range of potential X_{max} values would be prohibitive therefore we will quantify the deviations between our results and select Monte Carlo time series. Beating Monte Carlo methods for speed is a given once a reliable procedure for template construction has been established.

In terms of speed the immediate competition stems from fully analytical models. So far none of them succeded in matching Monte Carlo simulations even within expected measurement uncertainties therefore our approach needs to systematically outperform them in terms of accuracy.

First of all we should establish the baseline for our comparison. We simulated vertical proton- and iron-induced showers with three fixed primary energies of 10^{17} eV, 10^{18} eV and 10^{19} eV computing around 100 samples for each during the construction of our amplitude spectrum correction. Each of these showers can be used as both a template and a target real shower. This is not sufficient to cover all possibilities of "exotic" events but does reproduce the common distributions of shower maxima and other macroscopic quantities. There may be some innate bias induced by using the same set of simulations in both the amplitude spectrum fit and the test of its synthesis application when comparing average differences over many showers.

9.1 COREAS

As originally intended the primary benefit to our approach lies in the massively reduced computation time. Accurately determining the amplitude spectrum coefficient fits as per section 8.2 may require hundreds of showers with little numerical thinning at high energies but this would ideally be a one-time endeavour after which a published set of fitted values may be used. Similarly if completely unthinned showers are desired for use as templates the resulting radio time series can be saved at negligible cost. During end-user runtime our model need only load the pre-processed template(s) into memory once which may be a significant factor for older storage media but becomes negligible for large-scale production runs. A full synthesis of our 6 reference antenna positions completes in less than one minute including the load times, additional antennas can be added within seconds even for our completely unoptimised implementation as a single-thread python script. By contrast each real/template simulation takes over a week to complete with the single-thread implementation of CORSIKA.

Computing time would not appear to be a limiting factor given the rarity of UHECR air shower events but this low event rate necessitates extracting as much information as possible. Leading reconstruction techniques [5] require *hundreds* of dedicated Monte Carlo simulations *per* event which does make the computation demands a limiting factor in high-accuracy radio analysis even in existing dense detectors let alone any plans for coming generations.

As the status of analytical models shows rapid computation is not the primary hallmark of a good model however so we will now focus on exploring the accuracy of our model in various forms.

9.1.1 *Difference in Depth of Shower Maximum*

Starting with its observation in section 5.4 the systematic variation of our synthesis results with the difference in depth of shower maximum between template and target has been a driving force behind further development. As such we are very satisfied with our model now showing no further correlation in the peak amplitudes as demonstrated in figure 48 while still facing some inaccuracies for the most extreme differences when comparing the energy in the pulses as shown in figure 49. We also encounter some lingering numerical issues related to the stability of the amplitude spectrum fits and the resulting correlation parametrisations in the Charge Excess component and especially at the largest lateral distances. Using showers initiated by 10^{17} eV protons as templates as in these two figures is the worst-case scenario in this context but when omitting obvious malfunctions the peak amplitudes are still reproduced with around 15% deviation at most. Most lateral distances still show a lingering correlation between the deviation and ΔX_{max} as well as a a hierarchical ordering by lateral distance.

In terms of both stability and accuracy we have had significantly more success using showers initiated by 10^{19} eV iron nuclei as the templates for most choices of target. Due to their intrinsic statistics the horizontal axes in figures 50 and 51 differ from the previous plots but the eligible list of target showers was identical. Unfortunately the energy still deviates by up to 50% for Charge Excess emission and 25% for Geomagnetic emission for cer-

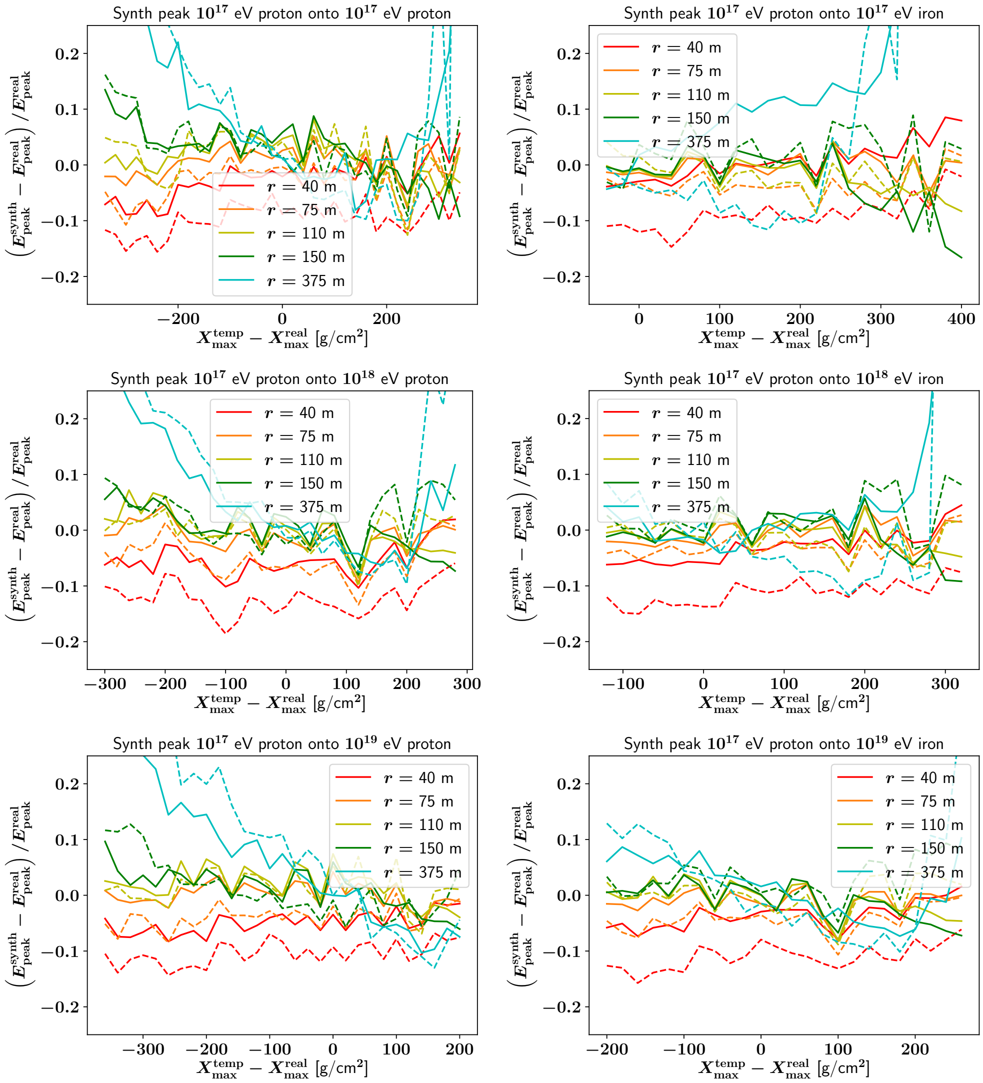

Figure 48: Relative deviation in peak amplitude between the synthesised and real time series for target showers initiated by protons (left) and iron nuclei (right) of 10^{17} eV (top), 10^{18} eV (middle) and 10^{19} eV (bottom) respectively using 10^{17} eV proton-induced showers as the templates. Each position on the horizontal axis represents a pair of showers with the stated difference in X_{max} and a selection of lateral distances. Geomagnetic emission is drawn with solid lines, Charge Excess emission with dashed lines. The large deviations for the largest lateral distance (375 m, cyan) are clearly identifiable as numerical malfunctions but we have no automatic means of correcting them.

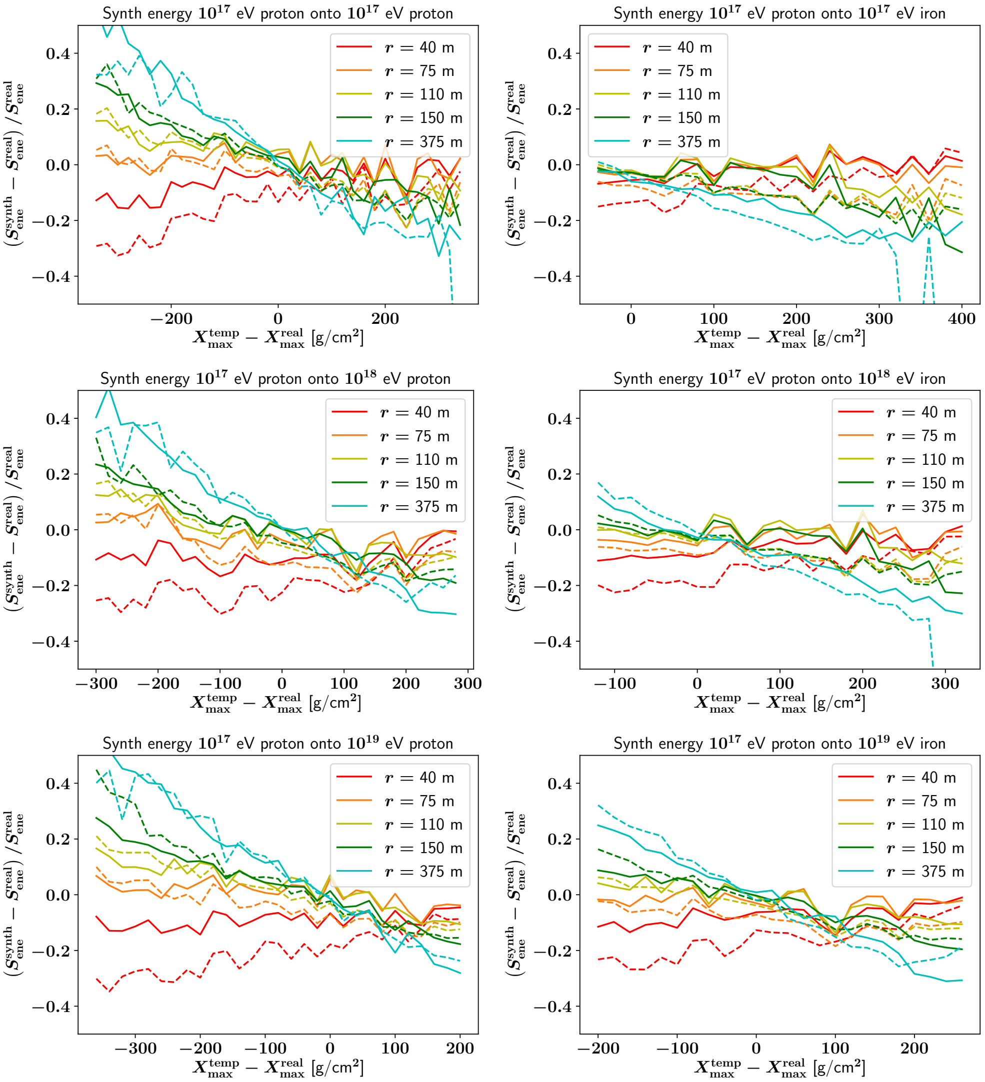

Figure 49: Same as figure 48 but now comparing the energy S_{ene} in the observable pulse approximated as the sum over the squared amplitudes in a window around the peak performed independently for each polarisation.

tain lateral distances and large differences in X_{max} for proton-initiated target showers. These deviations also appear to have a significant systematic bias towards overestimating the signal. Our ability to match proton- and iron-induced showers over a wide energy range represents a significant feat in its own right although the model configuration used here contains separate spectral fit coefficients for each primary species and energy.

9.1.2 *Shower-to-Shower Fluctuations*

In order to study the expected fluctuations under more typical conditions we pick a single 10^{19} eV iron-induced shower as our template. This specific shower was chosen for its X_{max} of 699 $\mathrm{g/cm^2}$ which is almost identical to the mean value for all 100 simulated 10^{19} eV iron-induced showers and close to the mean of all showers in our sample set at 683 $\mathrm{g/cm^2}$. All sample showers (except for the template itself) are then synthesised using this one template and the resulting differences averaged separately for each primary species and energy. This is equivalent to marginalising X_{max} as a nuisance parameter or sampling CORSIKA simulations with fixed physical inputs and no further pre-processing (CONEX can be used to pre-select X_{max} with a precision of $\approx$ 10 $\mathrm{g/cm^2}$ as mentioned in section 2.4).

The resulting peak amplitudes shown in figure 52 deviate by an average of no more than 6% for the Geomagnetic component (Geo) and 11% for Charge Excess emission (CE) which we interpret as the remaining systematic inaccuracy of our model and fluctuate around this mean by an additional 4% and 7% (1σ) respectively which constitutes the impact of shower-to-shower fluctuations when weighted with a "realistic" statistical distribution of X_{max}. As in the previous section the pulse energy matches less precisely at 15% (Geo) and 19% (CE) mean deviation as well as 11% (Geo) and 15% (CE) fluctuation both dominated by large lateral distances. Averaging with equal weights over all target primary species and energies as well as all lateral distances yields an accuracy of 2±2% for Geomagnetic emission and 4±2% for the Charge Excess component in terms of peak amplitude and for the pulse energy 4±3% and 8±4% respectively. Care should be taken not to read too much into these statistics in their relevance to real analyses however as the detailed plots as well as the previous section clearly show correlations between the level of deviation and X_{max}, lateral distance and the specific choice of templates.

Using the combined quadratic fit for the amplitude spectrum parameters increases the deviations to 5±5% (Geo) and 12±26% (CE) in peak height or 14±11% (Geo) and 17±21% (CE) in pulse energy at their largest. A comparison between figures 54 and 52 for the peak height or figures 55 and 53 clearly shows the reduced accuracy but also demonstrates that all sample

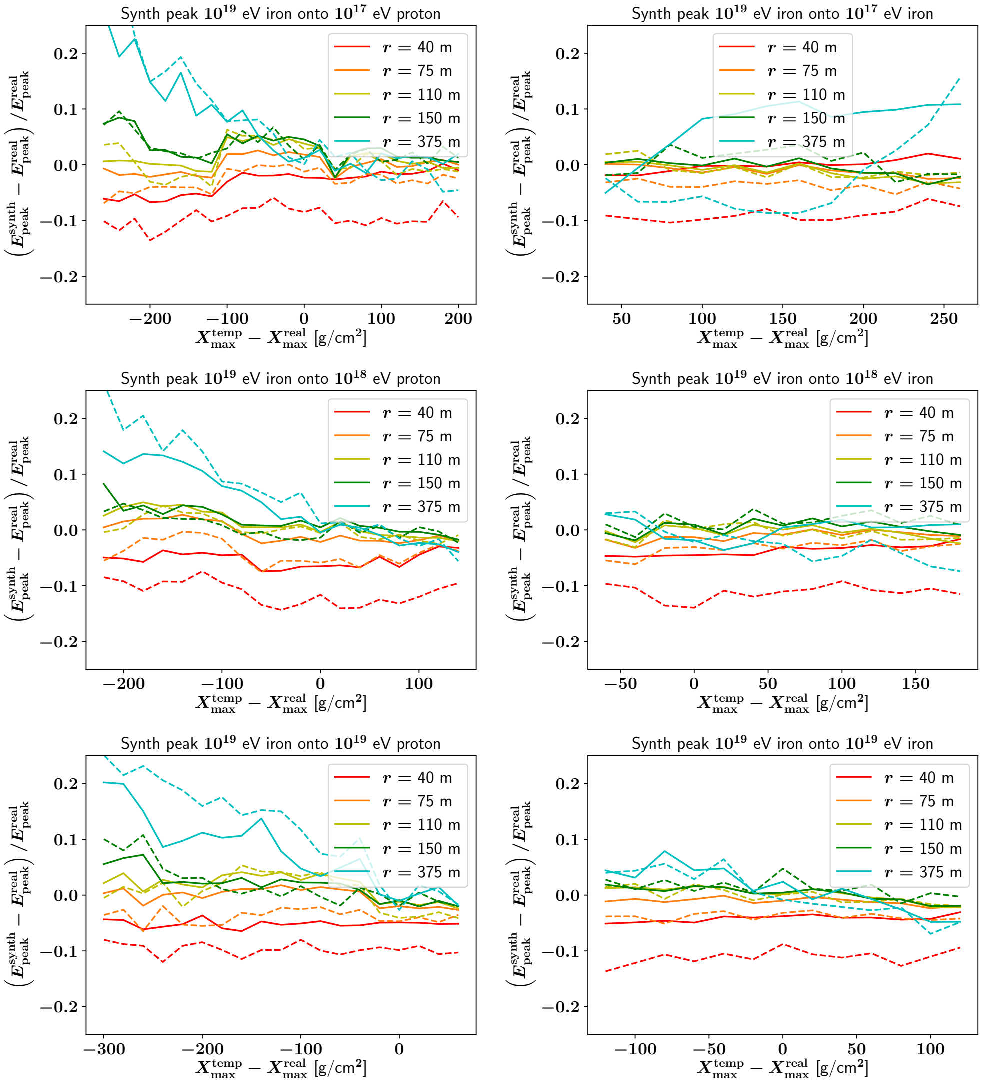

Figure 50: Same as figure 48 comparing the peak heights for various target populations now using showers induced by 10^{19} eV iron nuclei as the templates. Except for $r = 375$ m there appear to be no more variations with ΔX_{max}

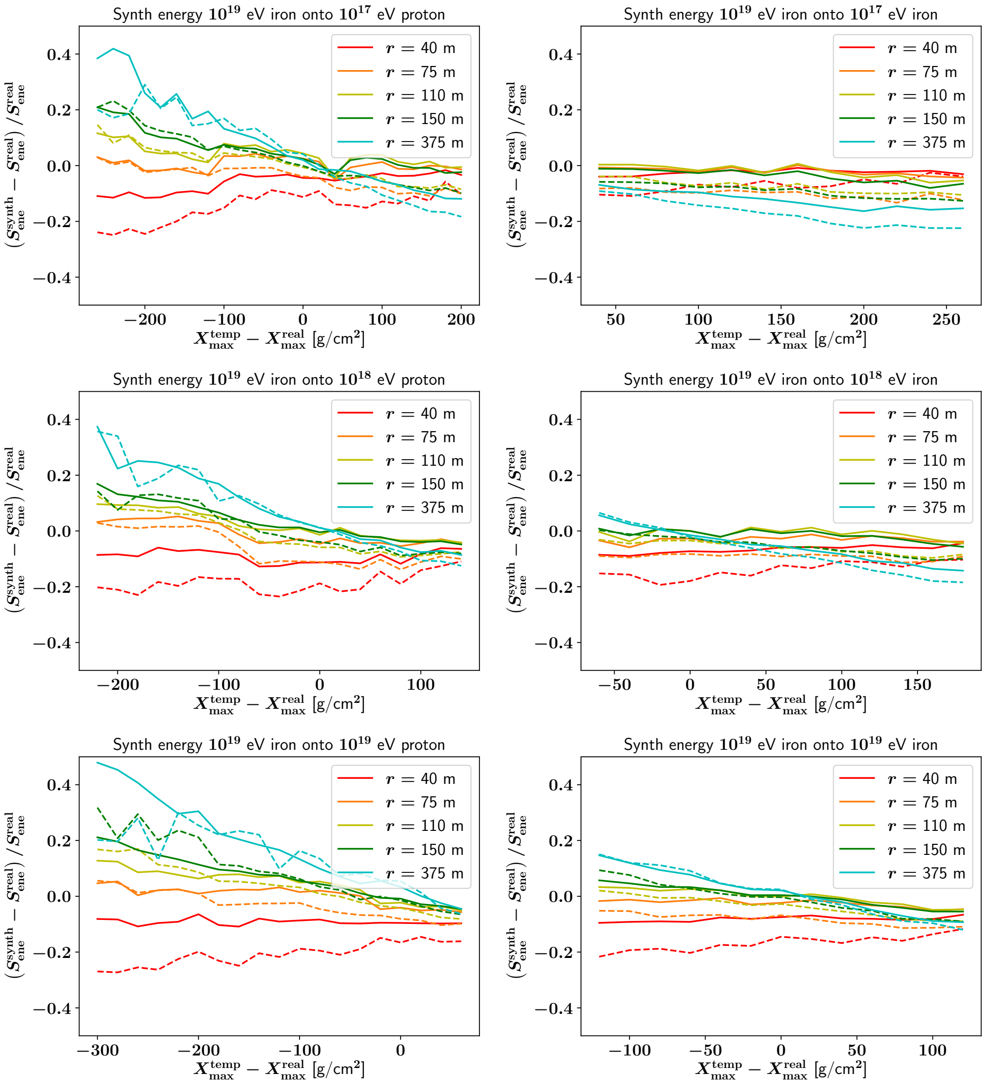

Figure 51: Same as figure 49 using 10^{19} eV iron showers as the templates but now comparing pulse energies.

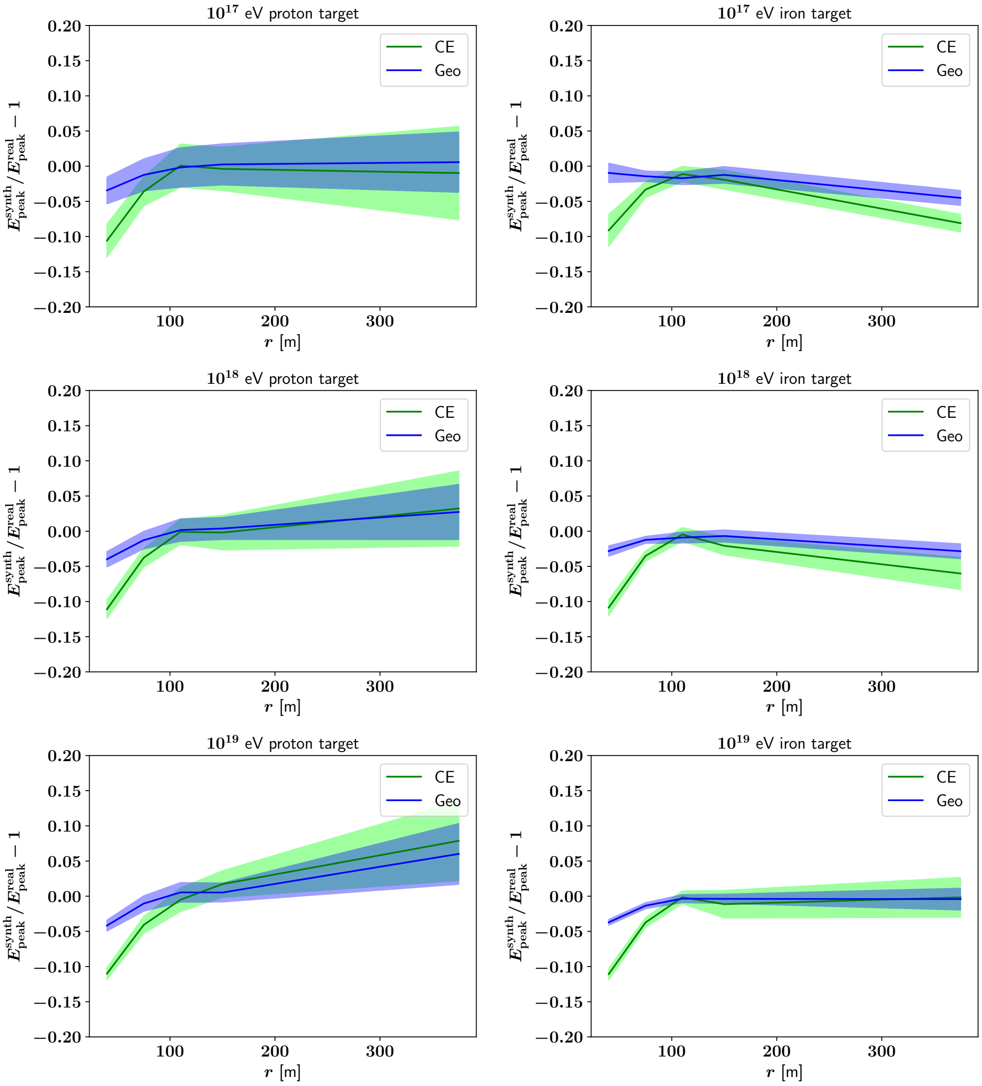

Figure 52: Average lines and 1σ deviation bands of the relative difference between the synthesised peak height and its true value for target showers initiated by protons (left) and iron nuclei (right) of 10^{17} eV (top), 10^{18} eV (middle) and 10^{19} eV (bottom) respectively using a single 10^{19} eV iron-induced shower with $X_{\mathrm{max}} = 699\ \mathrm{g/cm^2}$ as the template. Unlike the previous figures no distinction was made between the maxima of the target shower, the full simulation set was used meaning approximately 100 samples in each individual plot. Charge Excess emission is shown in green and the Geomagnetic component in blue.

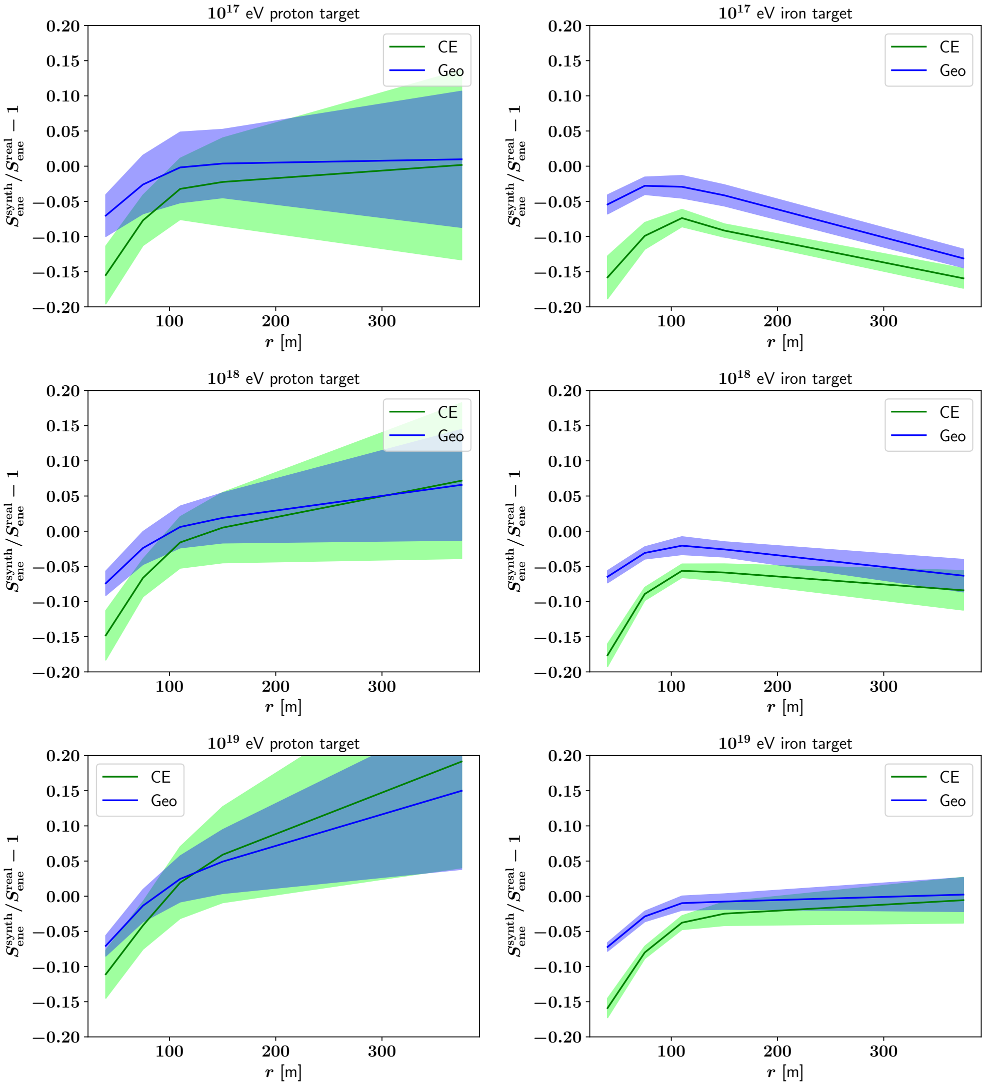

Figure 53: Average lines and 1σ deviation bands of the relative difference between the synthesised pulse energy in a window around the peak and its true value for target showers initiated by protons (left) and iron nuclei (right) of 10^{17} eV (top), 10^{18} eV (middle) and 10^{19} eV (bottom) respectively using a single 10^{19} eV iron-induced shower with $X_{\mathrm{max}} = 699$ g/cm^2 as the template. Unlike the previous figures no distinction was made between the maxima of the target shower, the full simulation set was used meaning approximately 100 samples in each individual plot. Charge Excess emission is shown in green and the Geomagnetic component in blue.

populations except the Charge Excess component of 10^{19} eV proton-induced showers are still synthesised with impressive accuracy.

9.1.3 *Time Series Comparison*

Generally the synthesised time series match the real signals almost perfectly with minor deviations primarily visible in the peak height at $r = 40$ m and the overall amplitude at $r = 375$ m as exemplified by figure 56. This represents a substantial improvement over all previous attempts with the one exception of our amplitude spectrum correction being able to drive signals at the largest lateral distances off their already near-perfect equality. This remains true across two orders of magnitude in primary cosmic ray energy and also when matching showers initiated by different species of primary particles. As an example figure 57 shows the synthesis of a 10^{19} eV proton shower using a 10^{17} eV iron-induced template with the resulting quality very similar to the previous example.

Using our pre-selected template serving as the basis for figure 52 and following also yields excellent results including the case where we fit the amplitude spectrum parameters across all primary energies and species simultaneously with a quadratic function. This configuration is closest to what we would recommend for large-scale production although it might be desirable to choose a lower template primary energy when computing unthinned templates as the computation time increases substantially compared to optimised thinning.

One remaining concern is the applicability and stability of the amplitude spectrum fit at large lateral distances as illustrated by the synthesis results in figure 59. Due to the low overall strength of the coherent signal incoherent contributions have a far greater influence compared to closer lateral positions which is most clearly demonstrated by the phase spectrum in figure 47 but also visible in the time series as high-frequency "wriggles". For vertical showers the primary failure mode occurs when the template has a particularly large value of X_{max} leading to very low statistics in earlier slices. As our amplitude spectrum fit is still performed over the full target frequency range it may begin to fit more and more noise leading to poor matching and volatility. Low template particle numbers exacerbate the issue as dividing by them for the longitudinal normalisation artificially inflates any numerical errors. As our current spectrum parametrisation is fitted to each lateral distance individually we could design a more robust fit function in the future. A more pragmatic option exists in simply disabling (or smoothly scaling down) the amplitude spectrum correction with increasing lateral distance past the Cherenkov ring as the simple synthesis model has worked almost perfectly for these distances from the beginning.

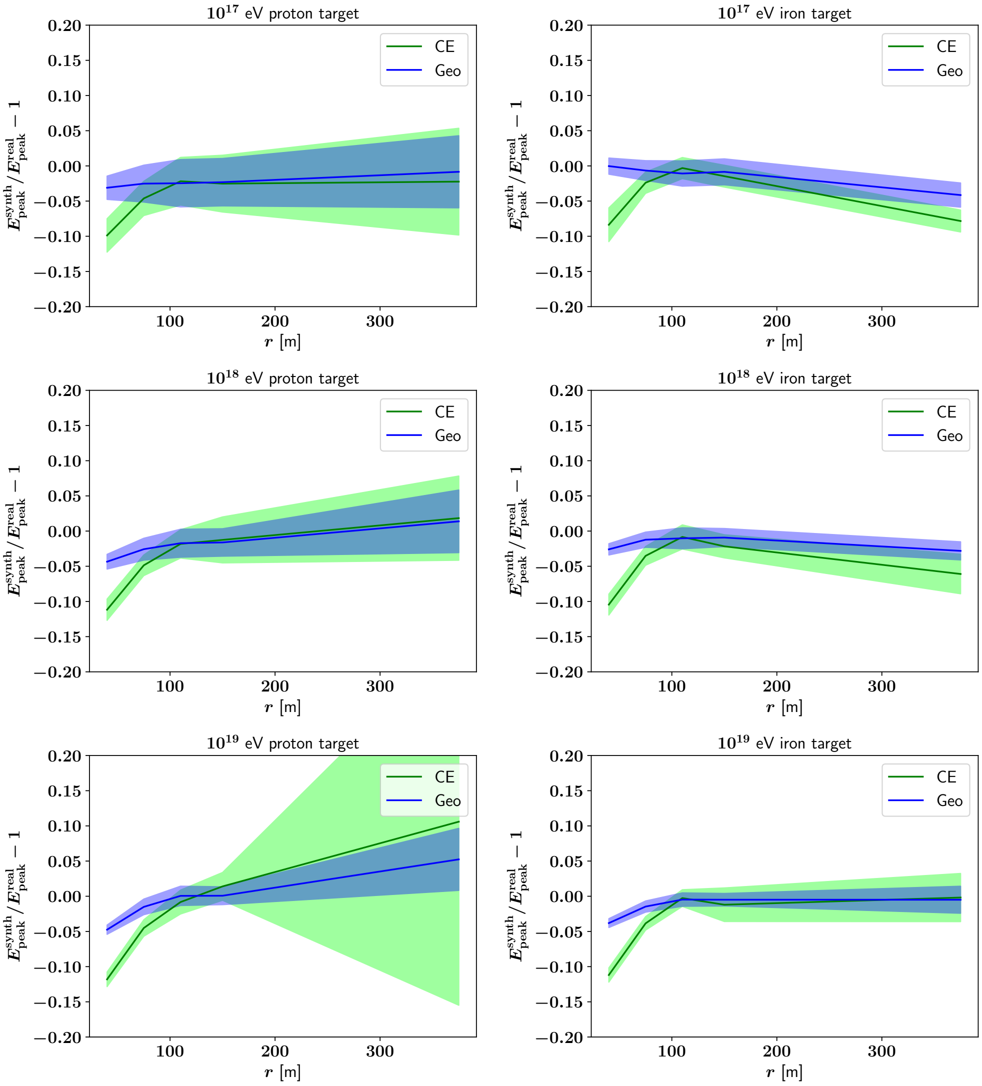

Figure 54: Identical setup to figure 52 but now comparing the relative deviation in peak amplitude when the synthesis is performed using a single amplitude spectrum correlation fit across all primary species and energies meaning the same coefficients were used in the creation of all six plots.

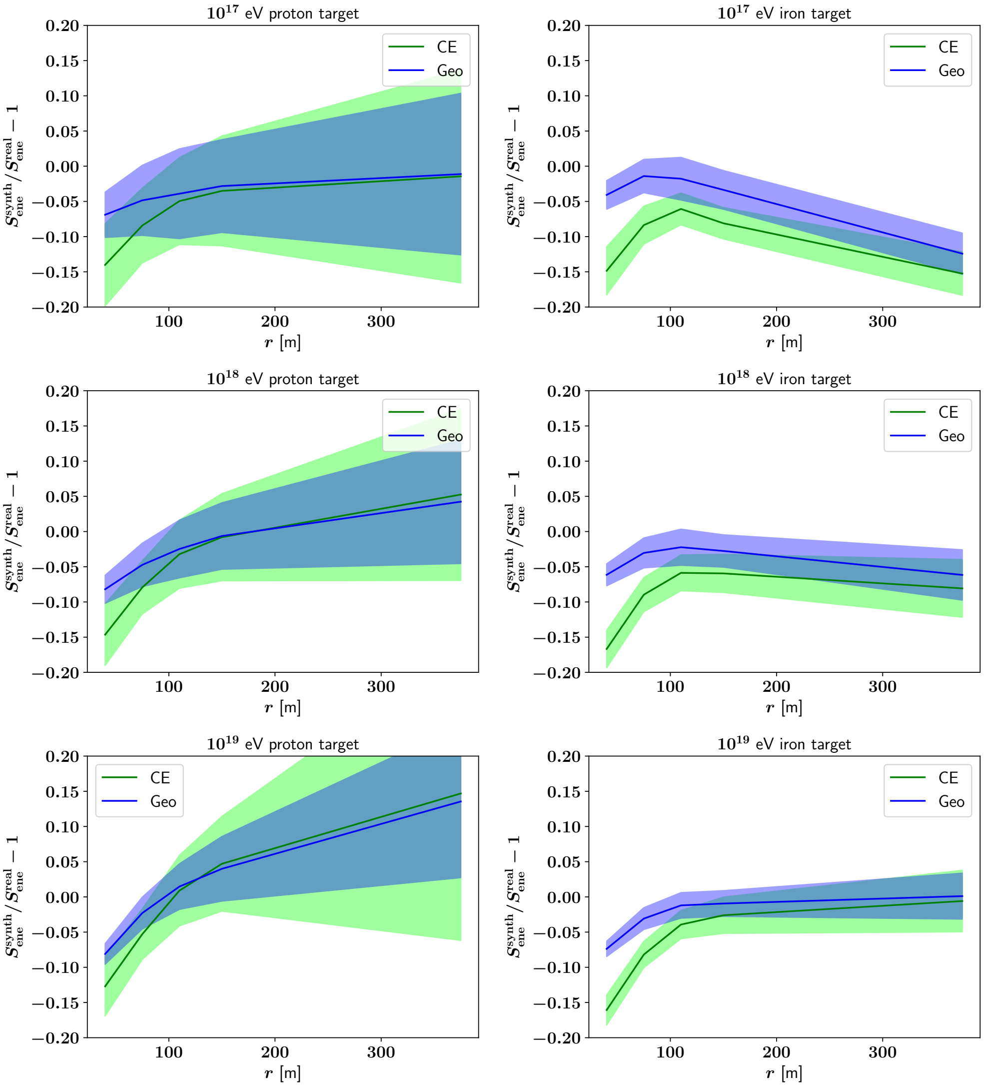

Figure 55: Identical setup to figure 53 now comparing the pulse energy in a window around the peak where the same coefficients were used for amplitude spectrum correction during synthesis for all six plots.

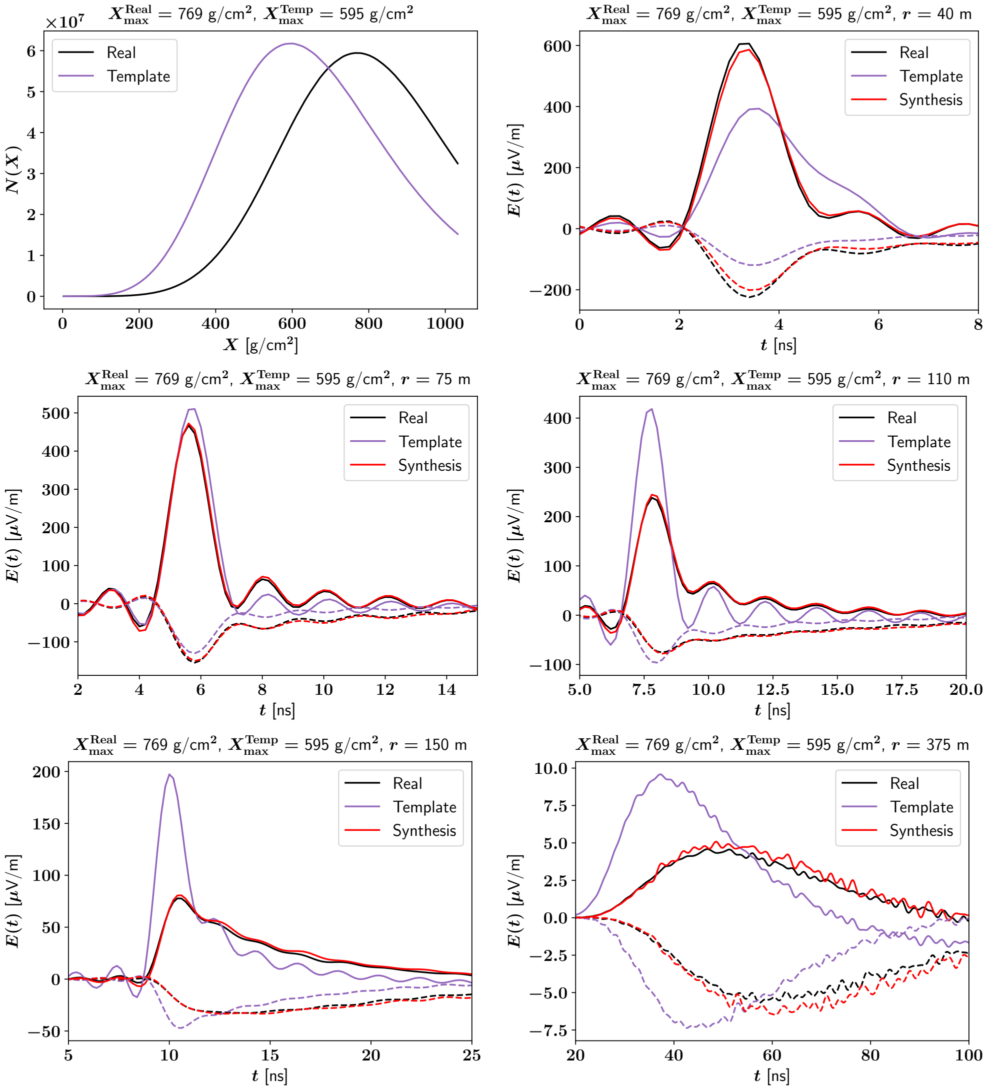

Figure 56: Full synthesis of a 10^{17} eV proton shower using a different 10^{17} eV proton shower with $\Delta X_{\max}$ = -174 g/cm^2 as the template. Despite the significant differences between both showers all lateral distances are synthesised almost perfectly for both components (solid: Geomagnetic, dashed: Charge Excess) with only minor deviations for the smallest and largest lateral distances.

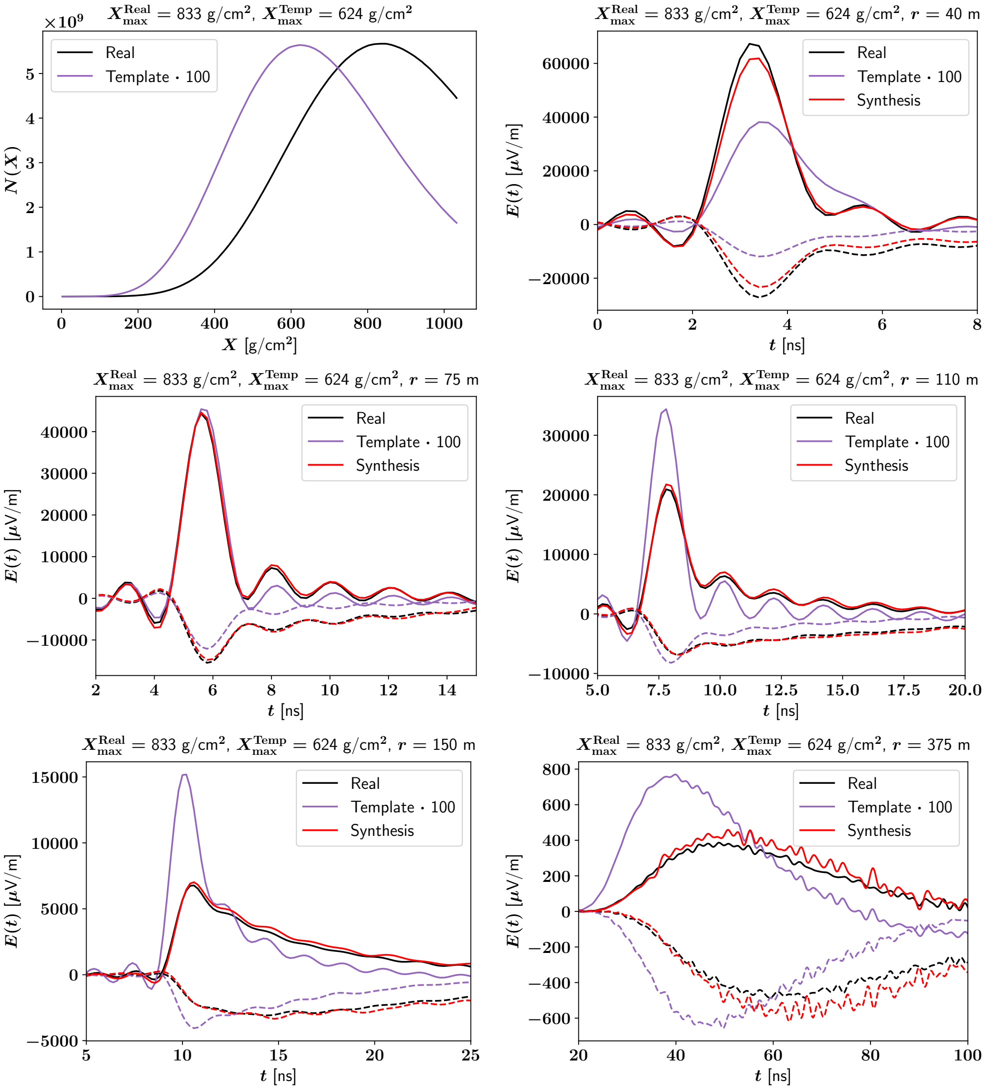

Figure 57: Full synthesis of a 10^{19} eV proton shower using a 10^{17} eV iron-induced shower with $\Delta X_{\max}$ = -209 g/cm^2 as the template. Our correlation coefficients for the amplitude spectrum fit were determined independently for different primary energies and species in this example which would require initial knowledge of these properties to reproduce in reality. Again all lateral distances are synthesised almost perfectly for both components (solid: Geomagnetic, dashed: Charge Excess) with only minor deviations for the smallest and largest lateral distances. Note the template shower was only scaled by a factor of 100 for display purposes not during synthesis, the effect of the primary cosmic ray energy was intrinsically encoded in the magnitude of the longitudinal profile.

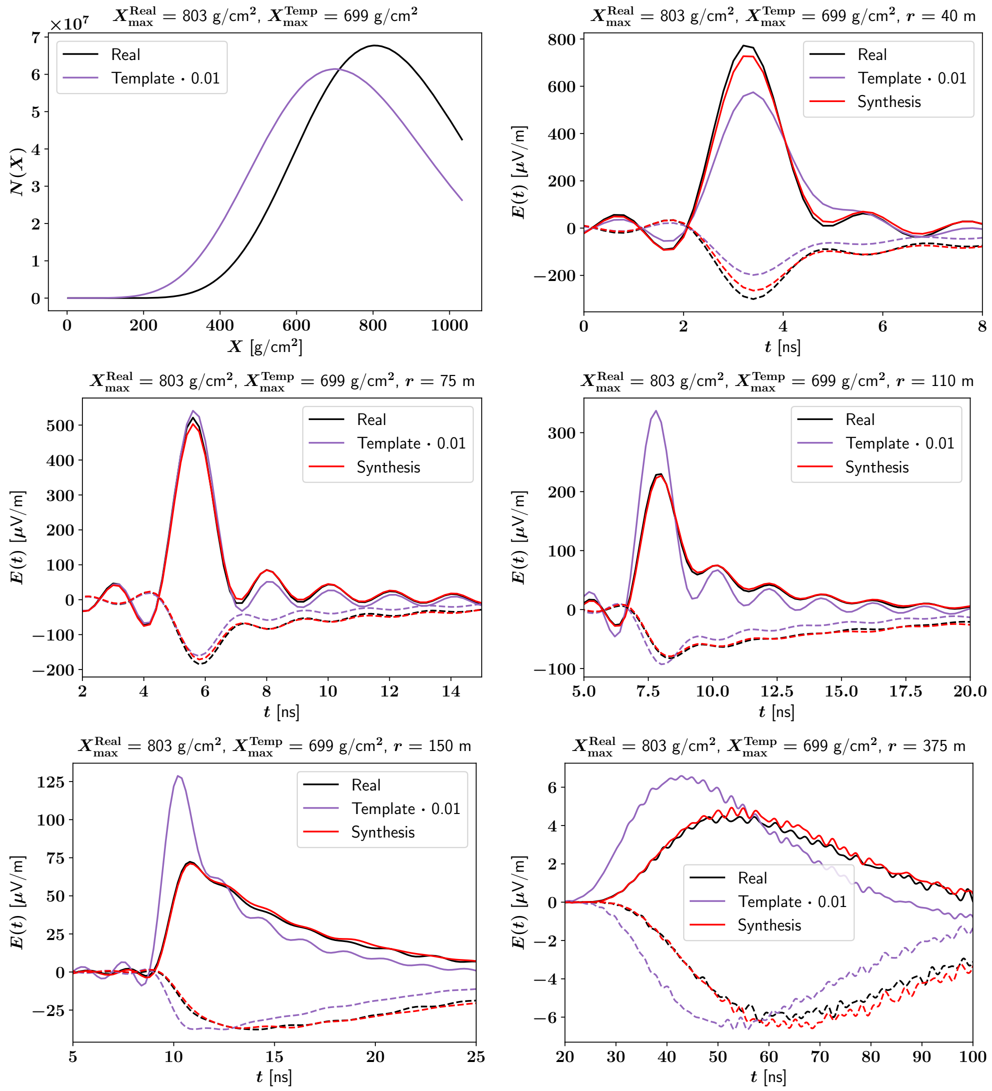

Figure 58: Full synthesis of a 10^{17} eV proton shower using our 10^{19} eV iron-induced shower from figure 52 with $\Delta X_{\max}$ = -104 g/cm^2 as the template. Here we use the combined quadratic correlation fit over all primary species and energies for both the template and the target shower meaning we do not assume any prior knowledge of the target primary particle beyond the longitudinal profile. Once again all lateral distances are synthesised almost perfectly for both components (solid: Geomagnetic, dashed: Charge Excess) with only minor deviations for the smallest and largest lateral distances proving the combined fit can also maintain the pulse shape. The template shower was only scaled by a factor of 0.01 for display purposes not during synthesis, the effect of the primary cosmic ray energy was intrinsically encoded in the magnitude of the longitudinal profile.

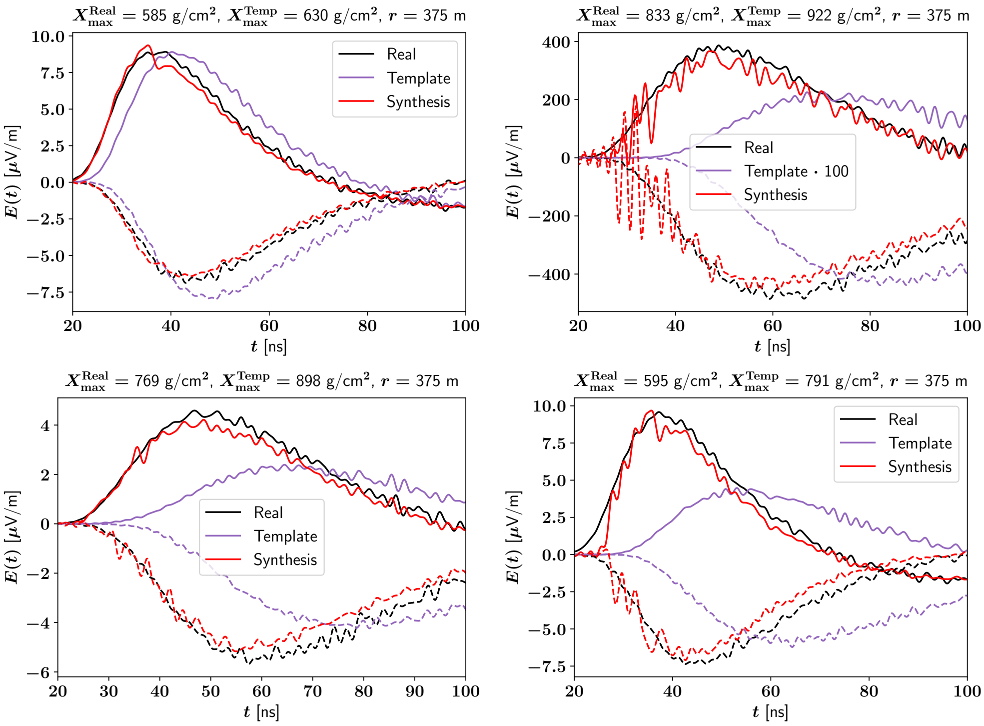

Figure 59: Examples of various failure modes for large lateral distances. All four examples use 10^{17} eV protons as the template and the bottom two also use (different) 10^{17} eV proton showers as targets. The top left target was induced by a 10^{17} eV iron nucleus, the top right target shower originated from a 10^{19} eV proton. In the top left a "kink" is visible in the Geomagnetic component synthesis result, likely a result of poor fit stability distorting a single slice. In the top right both polarisations show very obvious numerical artifacts in the rising slope of the pulse. Due to the extremely deep template X_{max} of 992 g/cm^2 coupled with the arrival time structure of the slices this represents a severe mismatch in the very early cascade development where the template slice radio signal may even be entirely incoherent while the target shower has already built up a substantial population of emitting particles. The bottom right plot shows a variation of this configuration where the fit converges to more standard values and the lack of radio signal dominates over the particle number rescaling leading to a hard cutoff. The bottom left plot shows a coherent mismatch in electric field amplitude proving showers far off the correlation line for our amplitude spectrum correction can in fact be distorted *away* from the near perfect synthesis of the simple particle number rescaling. In all pulses including the real and template signals high-frequency contributions are clearly visible and not fully reproduced in the synthesis result suggesting the valid frequency range for our synthesis may be limited to less than 500 MHz at large lateral distances.

9.2 MGMR3D

The MGMR3D model as introduced in section 3.1.3 constitutes our primary competitor in terms of computation speed. Due to being fully analytical such models are preferable to our template synthesis approach provided they can reach comparable accuracies. For computational operation MGMR3D takes the primary cosmic ray energy and the depth of the shower maximum X_{max} as its main physical inputs. The model features no explicit distinction between primary particles of different species but several secondary parameters could be tuned to better fit specific compositions. We opted not to do so and instead left all other parameters at the recommended default values. Inputs of potential physical interest include the Charge Excess fraction $J_Q^0 = 0.25$, the presumed depth of the first interaction $X_0 = 200$ g/cm^2 and the longitudinal scale parameter $\gamma = 100$ g/cm^2 ("LamX" in the input code). As MGMR3D includes no innate probabilistic components each of our Monte Carlo sample showers produces exactly one MGMR3D simulation. Samples at equal primary energy but induced by a different species of primary particle will yield identical MGMR3D outputs if their respective X_{max} values happen to be identical but such coincidences should be rare (and are not fulfilled exactly at our numerical X_{max} resolution however MGMR3D outputs are expected to vary smoothly with small differences in X_{max}).

9.2.1 *Scale Calibration*

One significant shortcoming of MGMR3D in this regard is the lack of an absolute scale calibration at least at the time of writing. This scale can be estimated to a reasonable degree by comparison with CoREAS or multi-component reconstructions of real air showers but no canonical value has been published. We therefore somewhat arbitrarily match the averages over showers of all primary species and energies in our sample set considering only the Geomagnetic emission component as the Charge Excess fraction is a free input parameter in the code (though we did not vary it in the context of this work). As MGMR3D is known to be less accurate for large lateral distances [41] we pick a fixed reference distance of 110 m instead of computing the lateral average. Different choices of the reference distance may change certain results as by construction the mismatch will be smallest directly at this distance. Averaging the peak ratios for each set of showers with equal primary particles yields the scale conversion values shown in table 1 which average to one MGMR3D numerical electric field unit equaling 3.2 nV/m. This single value is used globally in all our comparisons. While the values for proton primaries are consistent within their fluctuations iron-induced showers instead show a systematic energy dependence while shower-to-shower variations at fixed en-

Table 1: Empirical absolute scale calibration of MGMR3D numerical radio units.

in nV/m	10^{17} eV	10^{18} eV	10^{19} eV
proton	3.2±0.3	3.4±0.2	3.5±0.2
iron	2.6±0.07	3.0±0.06	3.3±0.04

ergy remain small. Thus there appears to be a systematic deviation which may relate to a changing slope of the lateral distribution. As the MGMR3D model has no explicit representation of the primary cosmic ray species any number of second-order cascade effects not represented in X_{max} could be responsible. Investigating these details goes beyond the scope of this work and would imply further refinement of the MGMR3D model beyond its current implementation.

9.2.2 *Shower-to-Shower Fluctuations*

Figures 60 and 61 were constructed using the same layout as those in section 9.1.2 and thus should be eligible for direct comparison. Immediately one must note the significantly larger vertical scale, our synthesis model plots are clamped to 20% variation for visibility while MGMR3D deviated by as much as 100% in the peak height and even 500% in the Charge Excess pulse energy at the largest lateral distances and also 100% for the Geomagnetic component.

The average deviations are 26±6% (Geo) and 42±6% (CE) in peak height or 77±12% (Geo) and 180±18% in pulse energy meaning our synthesis model is several times more accurate in simulating "realistic" X_{max} populations of cosmic ray primaries across two orders of magnitude in energy and most likely spanning the full range of primary masses *in addition* to having an innate absolute scale calibration.

9.2.3 *Time Series Comparison*

Correctly synthesising the pulse shape is where our model truly shines. Figures 62, 63 and 64 show repeats of the previous figures 56, 57, 58 this time with the MGMR3D simulation of the real signal added in. The peak amplitude can be matched quite well up to lateral distances of 150 m for Geomagnetic emission (after calibrating the global scale at 110 m) and 110 m for Charge Excess emission but the pulse shapes simply do not align well. The pronounced "ringing" in the real and synthesised signals is an artifact of our bandpass filter but drawing an envelope function by eye still reveals significant differ-

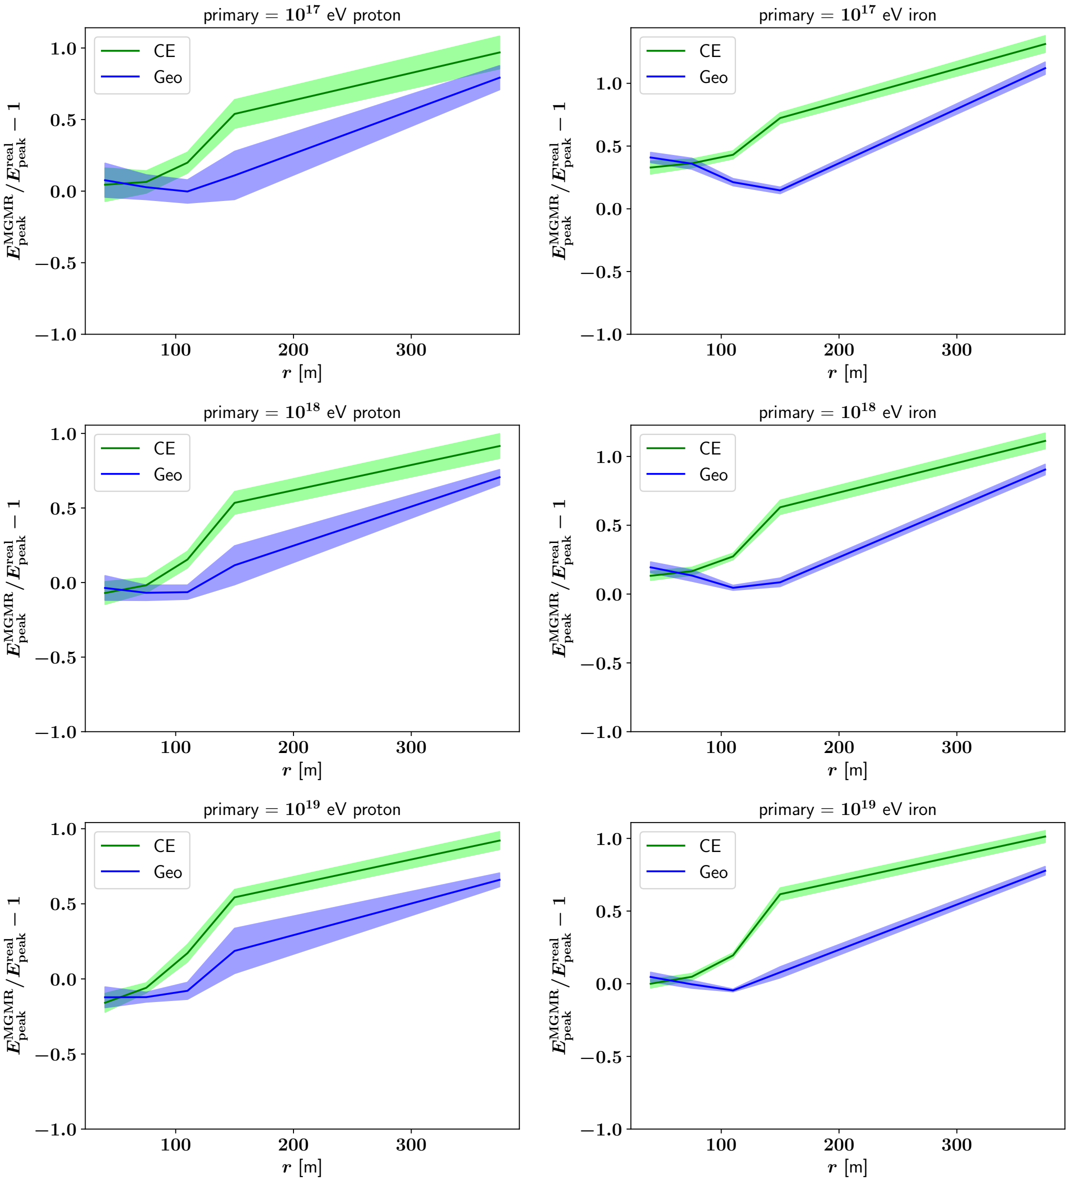

Figure 60: Mean and standard deviation (1σ bands) of the relative difference in peak height between MGMR3D simulated time series and the corresponding CoREAS results for target showers initiated by protons (left) and iron nuclei (right) of 10^{17} eV (top), 10^{18} eV (middle) and 10^{19} eV (bottom) respectively. The full simulation set was used meaning approximately 100 samples enter each individual plot. The fluctuations originate entirely from variations in the real radio signal and $X_{\max}$ which is the only changing input for individual samples of fixed primary species and energy. The absolute amplitude scale was tuned to the Geomagnetic peak height at $r = 110$ m therefore its average deviation should be zero, numerically we obtain a mean of 1.1% due to rounding of the stated value.

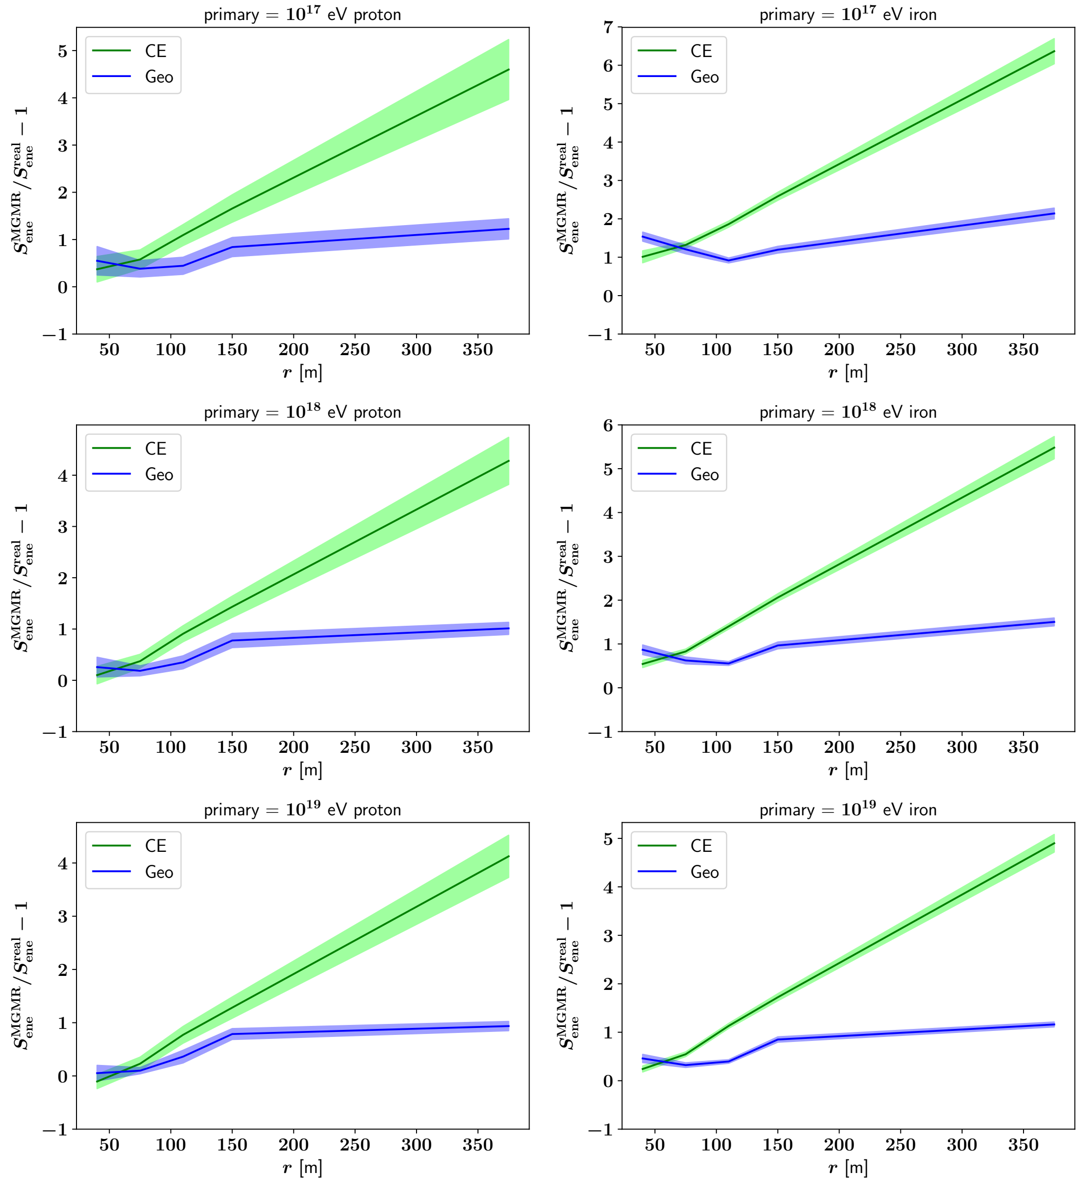

Figure 61: Same as figure 60 but now comparing the mean and standard deviation (1σ bands) of the relative difference in pulse energy within a window around the peak. Unlike the peak amplitude the energy was *not* tuned directly hence its deviation at $r = 110$ m remains valid.

ences in the slope even for smaller lateral distances. At large lateral distances MGMR3D can match neither the template nor the target shower despite being given the true Monte Carlo X_{max}.

Figure 65 shows a repeat of figure 59 demonstrating the innate differences in inaccuracies between the two models. Whereas MGMR3D fails to accurately reproduce the peak height and pulse shape at the bimodal level it inherently produces smooth time series devoid of incoherent contributions. Our synthesis model meanwhile, despite clearly containing numerical artifacts and errors, still reproduces the peak height with impressive accuracy and if one were to manually smooth the time series (e.g. via convolution with a window function of 5-10 ns) the remaining pulse shape would match quite closely as well.

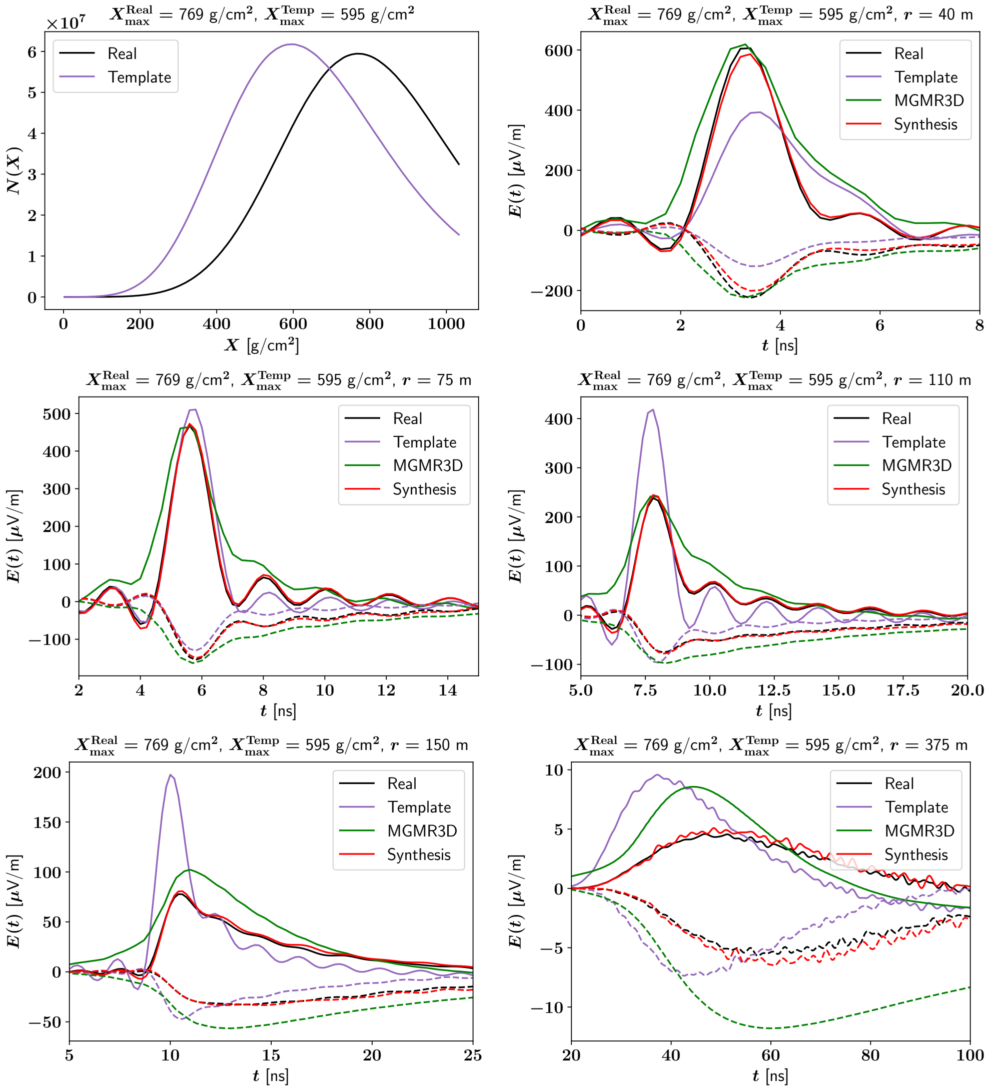

Figure 62: Repeat of figure 56 with the MGMR3D simulation of the target shower added in green. As MGMR3D contains its own frequency band settings its output was not filtered again which may explain the lack of "ringing".

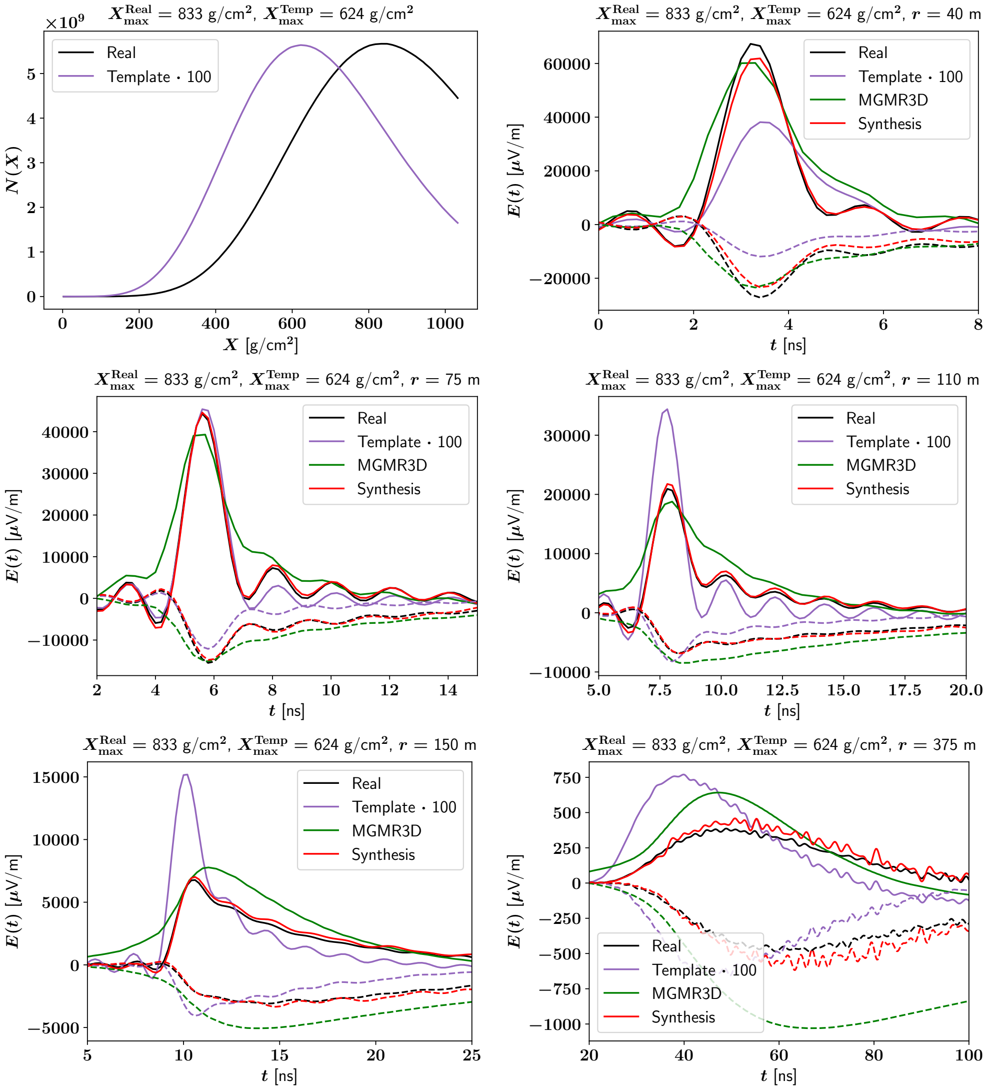

Figure 63: Repeat of figure 57 with the MGMR3D simulation of the target shower added in green.

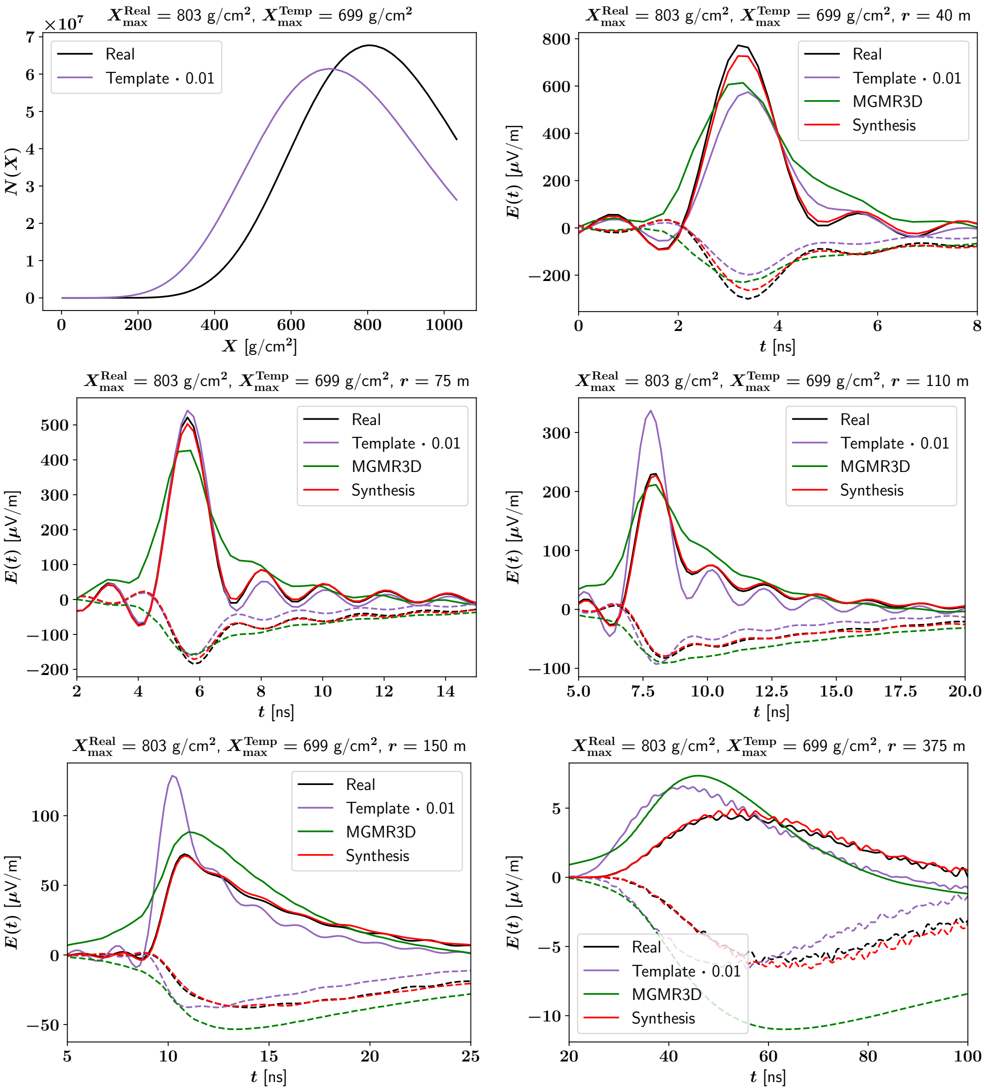

Figure 64: Repeat of figure 58 with the MGMR3D simulation of the target shower added in green.

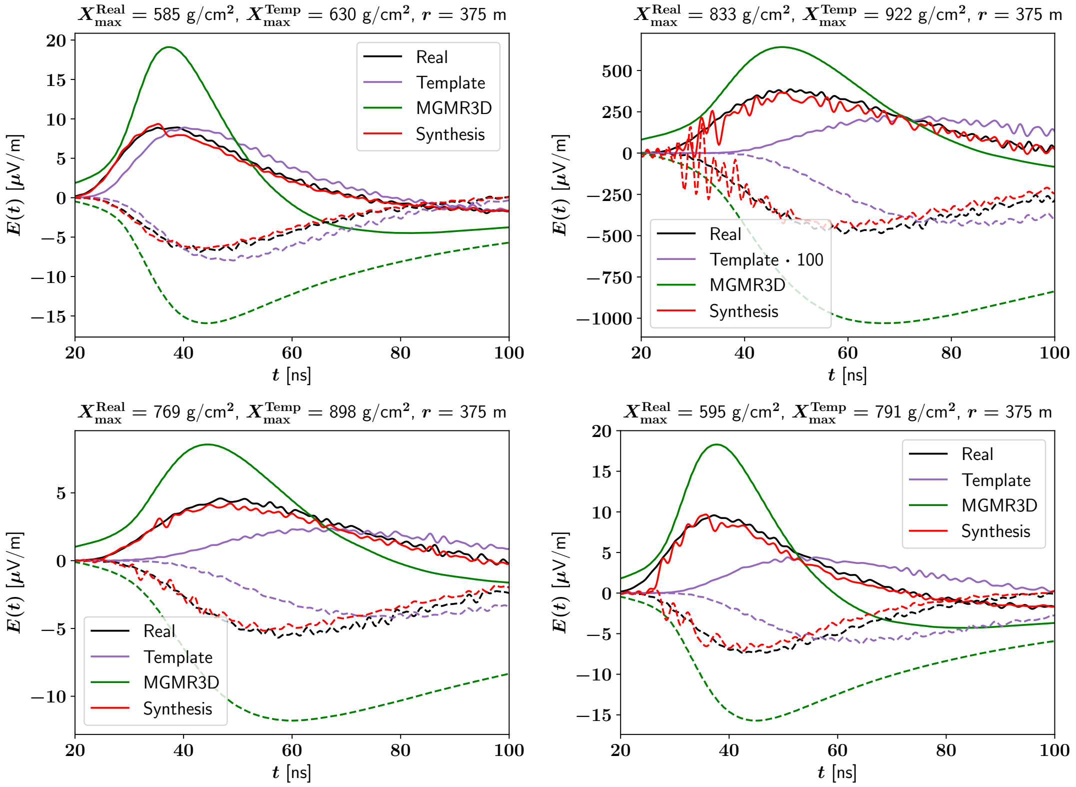

Figure 65: Repeat of figure 59 with the addition of the MGMR simulation results for the target shower in green.

10

FUTURE AVENUES

10.1 THEORETICAL STUDY OF THE AMPLITUDE CORRECTION FACTOR

So far our amplitude rescaling model works very well but the correlation between X_{max} and the radio amplitudes was determined empirically. In the future we would like to identify the physical relations giving rise to this correlation both to further our theoretical understanding of radio emission from cosmic ray air showers and for further refinements of our model. If an entirely analytical description could be found it might enable a new MGMR-based model to match the accuracy of our rescaling without using any templates. In the long run one could even design a next-generation probabilistic code capable of matching the output statistics of CoREAS without performing a detailed Monte Carlo simulation.

Regarding the amplitude correction factor specifically we know it represents a frequency-dependent adjustment of the radio output from a slice. The current model features both longitudinal and lateral variations which implies a more complex origin than a second scalar longitudinal component. We currently believe the energy distribution of the electrons and positrons as discussed in section 4.1 may be the determining factor as is known to evolve with shower age. An aging energy distribution would also lead to variations in lateral spread and particle velocities which tie to the geometric scale of radio coherence and may provide a more direct explanation of the lateral/angular dependence of the radio signal.

More theoretical groundwork is necessary however in order to properly understand our parametric and simulation-derived model, for example it should be possible to derive the pulse emitted by a source which is not merely a point moving along the shower axis but an extended distribution of particles (or charge/current density) which can then be compared to slicing configurations such as those detailed in section 6.6 in order to confirm which aspects of the cascade particle ensembles are truly relevant for radio emission.

10.2 STATISTICAL GENERALISATION

In this work we have only discussed the synthesis based on individual templates. Lacking a completely analytical understanding one could still compare and combine multiple template showers in an attempt to match the real shower-to-shower fluctuations present in CoREAS. Care must be taken with end-to-end machine learning approaches such as neural networks or Bayesian inference as they might learn unphysical inadequacies in the model rather than the physical aspects which are represented propery. Bayesian approaches such as those based on Information Field Theory [7] have a certain advantage here as their formulation itself provides a means of verification and direct comparison to statistical distributions extracted directly from CoREAS.

10.3 PRIMARY GEOMETRY (θ, φ)

On the other side we also need to encompass all quantities of interest to cosmic ray astronomy in order to make our model viable for application in real experiments. The arrival direction of the cosmic ray represents the greatest remaining challenge. With increasing zenith angle θ the shower develops further away from the antenna which does benefit us in regards to potential near-field effects. However at smaller zenith angles the electromagnetic component still impacts ground at sea level while reaching later development stages not present in our vertical template showers. For initial testing one can simulate dedicated templates for additional zenith angles but in the long run a continuous generalisation is required.

In very inclined showers, typically meaning zenith angles above 60° in experiments [50], the specific alignment of our slice borders may become relevant as layers defined orthogonal to the shower axis will include vastly different ambient densities while layers of constant slant depth may cover different evolution stages of the cascade above and below the shower axis. The curvature of the atmosphere may induce additional asymmetries in the propagation physics above and below the shower axis as already observed in experiments.

The azimuth angle φ is far more relevant in the context of measurements as opposed to simulations and modeling because its primary influence on the radio signal stems from alignment with respect to the magnetic field which we will discuss in the following. For very inclined showers variations in the azimuth angle will likely amplify most anticipated discrepancies.

10.4 MAGNETIC FIELD ($\vec{B}$, α)

To first order the magnetic field should only affect the Geomagnetic emission component. Per classical theory the electric field amplitude produced by the Geomagnetic mechanism should be directly proportional to the magnetic field strength and the *geomagnetic angle* α defined as the angle between the magnetic field and the shower axis

$$\vec{E}_{\text{geo}} \sim B \cdot \sin(\alpha) \tag{37}$$

with

$$\hat{v} \cdot \hat{B} = \cos(\alpha). \tag{38}$$

CORSIKA assumes the Earth's magnetic field to be homogeneous therefore this constitutes a scalar constant which does not depend on any antenna positions or source slicing. Empirical simulation studies suggest the emission strength might scale with an exponent closer to 0.9 than to unity [24] for the geomagnetic field amplitude while confirming the direct proportionality to the geomagnetic sine.

Deflection of cascade particles by the geomagnetic field induces minor asymmetries but since the original arrival direction is not known in measurements these will not be observable without a very dense next-generation hybrid detector. As the radio technique mostly sees use in the context of the highest cosmic ray energies and corresponding very low fluxes there is no broader community incentive to build such a detector in practice and a radio-only primary array cannot sample the particle footprint accurately enough to establish a correlation. On its own the prospect of studying radio emission from individual air showers in detail so far has not justified the financial and political investments necessary for such an experiment.

One rare but theoretically observable edge case of note occurs when the shower axis and the magnetic field are almost parallel. Due to the relative strengths of the emission mechanisms this immediately reduces the overall measured amplitudes by a factor of 5-10 but these rare events could offer us an opportunity to measure almost pure Charge Excess emission (plus the physically real incoherent contributions). In a more general sense small geomagnetic angles necessitate an accurate treatment of the Charge Excess emission in addition to the otherwise dominant Geomagnetic component whenever such showers can be detected above the noise threshold of the antenna system.

10.5 RADIAL SYMMETRY

Our synthesis core so far only covers discrete antenna positions along the $\vec{v} \times \vec{v} \times \vec{B}$ axis. In order to be of use in practical applications this needs to be generalised to the full lateral plane or even arbitrary three-dimensional antenna positions. Fortunately interpolation and extrapolation on the level of physically observable time series is routinely performed in existing analysis techniques [24] therefore our model can benefit from new developments in the broader community. The basic 2D computation outlined here is comparable to MGMR and has not been extensively tested with our synthesis results.

A key reduction of phase space arises from the expected radial symmetry of the particle ensemble and thus the radio signal. This symmetry is broken by the geomagnetic field if particles are able to drift especially since the cascade

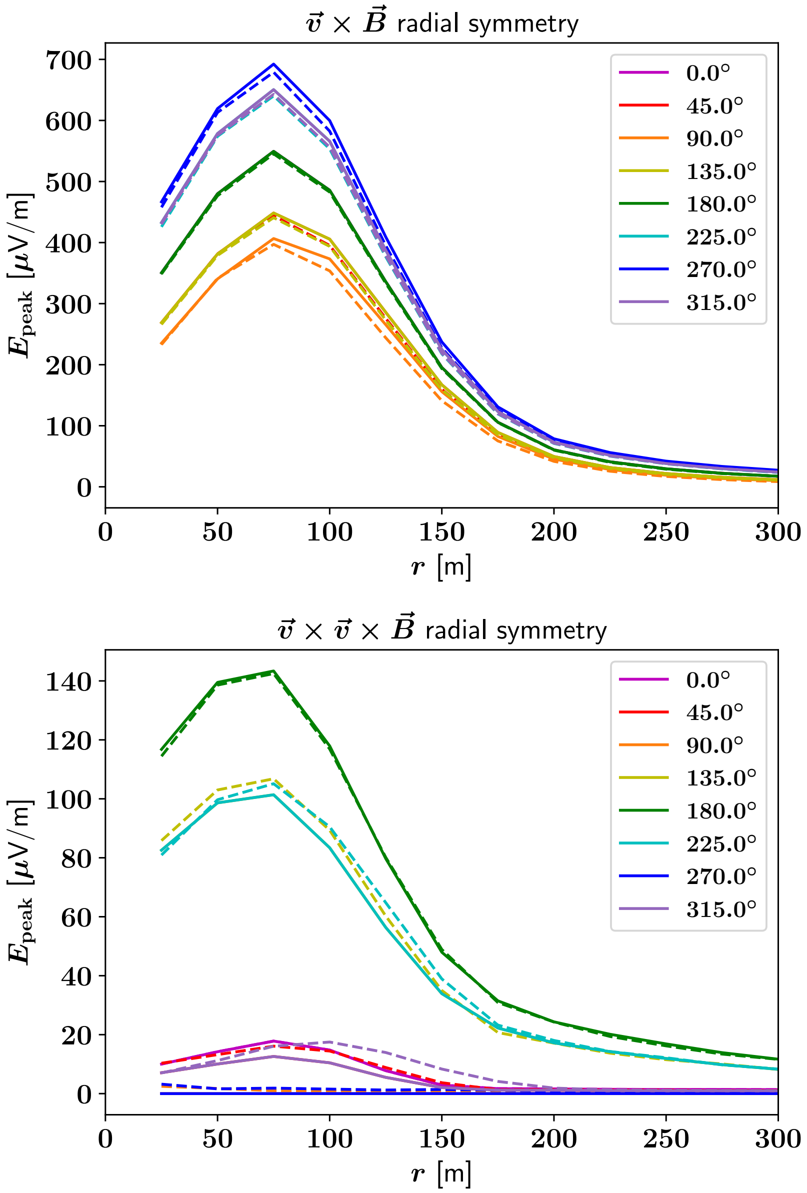

Figure 66: Test of radial symmetry for the $\vec{v} \times \vec{B}$ (top) and $\vec{v} \times \vec{v} \times \vec{B}$ (bottom) polarisation of the radio signal. Each curve shows a straight line of antennas with the stated angle relative to the $-\vec{v} \times \vec{v} \times \vec{B}$ (CoREAS $+\hat{x}$) direction once simulated directly in CoREAS (dashed lines) and once constructed from the 0° line based on our assumption of the two classical emission mechanisms decoupling for 0° and 180°. The $\vec{v} \times \vec{v} \times \vec{B}$ polarisation should only contain Charge Excess emission (whose projection should also drop to 0 at 90° and 270°) while the $\vec{v} \times \vec{B}$ polarisation may contain both Geomagnetic and Charge Excess contributions.

contains more electrons than positrons. We currently assume the amplitude of radio emission remains symmetrical for each of the two classical emission mechanisms individually. Thus the observed asymmetry of the radio footprint arises solely from interference between the two components allowing an analytical computation of the electric field as a vector sum as shown in figure 66. More complex models could also be adapted if necessary provided the polarisation-resolved time series along a single lateral axis suffice as inputs similar to the core of MGMR.

In inclined showers, in this case for zenith angles above 65° [51], additional asymmetries are known to arise due to the lines of sight traversing different paths in the inhomogenious atmosphere. The effects outlined and corrected for in existing publications should mostly be accounted for by properly considering the longitudinal extent of the air shower in our model. However most descriptions only consider the energy fluence while we synthesise the full time series.

10.6 LATERAL INTERPOLATION

Currently there are two interpolation methods under consideration in order to reduce the required number of lateral positions in our template library. As the long-distance geometrical correlations of the radio footprint are not well understood on the time series level we limit ourselves to linear interpolation between two template antennas. The main choice to be made is the domain in which the interpolation is calculated. Operating in the time domain allows for a numerically simple and robust implementation. On the theoretical side a major advantage is the linearity of this approach which greatly simplifies statistical analyses of the model. It does however carry a risk of amplifying spikes from incoherence or numerical noise.

Interpolating in the Fourier domain avoids this issue as the noise is isolated in high-frequency components. Indeed an independent linear interpolation of the Fourier amplitudes and phases has proven superior to the time-domain approach for our intended context as shown in [52], an example of the achievable accuracy is shown in figure 67. Unfortunately the computation of the phases is no longer linear which may pose a challenge in the future. The main drawback however is the risk of numerical errors. While rounding and minor inaccuracies are a necessary evil inherent to all computations the phase spectrum is prone to rare but severe failures: erroneously interpolating across a 2π periodicity bound for example can result in a phase shift of π equivalent to a sign flip of the entire resulting time series.

A potential safeguard against some common issues would be to overlay the highest peaks before performing the interpolation then substituting the proper arrival time independently. In the time domain such an approach

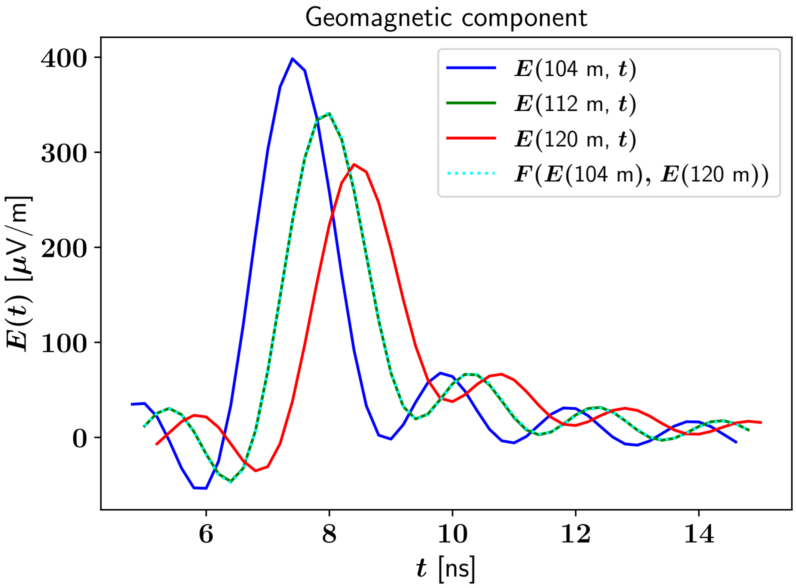

Figure 67: Plot of radial synthesis of the Geomagnetic component within a single air shower. The amplitude and phase spectra were interpolated linearly.

obviously requires the curve shapes to be reasonably similar especially in regards to zero crossings. In the frequency domain zeroing out the arrival times would greatly reduce the slope magnitude of the phase spectrum and thus leave far fewer crossings of periodicity bounds. Both variants require a reliable analytical timing model in order to recover the proper arrival time at the target position.

11

SUMMARY

We have successfully contructed a template synthesis model which can accurately reproduce results from full CoREAS Monte Carlo-based simulations. As it is intended for massive repetition in the context of inference or high-density next-generation radio antenna arrays [6] the initial load time of the templates should become negligible while the synthesis computation itself operates at speeds comparable to analytical models and orders of magnitude faster than Monte Carlo codes. Key aspects of the radio signal are matched with an impressive accuracy of better than 6% (Geomagnetic) and 11% (Charge Excess) in the height of the time series peak and better than 15% (Geomagnetic) and 19% (Charge Excess) for the pulse energy in a window around the peak for the remaining systematic deviations which may be possible to eliminate through further refinements. The quality of our results also includes a statistical component originating from shower-to-shower fluctuations on the order of 4% (Geomagnetic) and 7% (Charge Excess) for the peak height or 11% (Geomagnetic) and 15% (Charge Excess) for the pulse energy. These values approach the level of known limits on universality of the emitting particle cascades and thus may remain an irreducible limitation to our template synthesis approach. Further improvement in this regard may require adding a probabilistic component such as utilising neural networks or Information Field Theory [7] to learn the probability distribution of the shower-to-shower fluctuations and draw realisations from it as an addition to the analytical template synthesis.

Barring some numerical issues which could largely be solved through a more robust code implementation and careful vetting of the template showers the overall pulse shape is also reproduced remarkably well. As most systematics scale with the difference in X_{max} between template and target it may be useful to determine a broad estimate of the expected values of X_{max} during an actual analysis and to select templates accordingly. One potential solution to limited longitudinal coverage in a single template would be to select slices from multiple showers each close to their respective maxima thus guaranteeing high statistics across the full longitudinal range. Due to the importance of interference between subsequent slices a means of smoothing the transition may be required lest it introduce unphysical artifacts through the breaking of correlations in the template slices.

On the theoretical side our study of radio emission from individual slices has broadened our understanding of radio emission from air showers. On one hand empirically derived methods motivated by simple point source approximations enjoy a great deal of success in analyses of real data, on the other the longitudinal interference we observe between slice contributions and the complications it has caused would suggest these methods could never work. The solution likely lies in the context of the application: whereas we studied the full time series over a wide frequency range current analyses are limited by narrow measurement bands (typically 30-80 MHz for current experiments) and focus on calorimetry and stability with respect to noise. Most of the effects observed in our slice synthesis redistribute the radio emission to different viewing angles or arrival times and would thus continue to be registered *somewhere* in the data when integrated over the lateral plane or time axis respectively. The high degree of universality observed even at the level of individual slices also validates end-to-end correlation methods relating radio features to universality-defining parameters such as X_{max}.

Our most important open question remains the physical origin of the amplitude spectrum correlation we determined empirically from a large set of samples. We believe the shape of the spectrum ought to relate to the energy distribution of the electrons and positrons within the emitting slice which may also partially manifest in their lateral distribution. Through modification of the CoREAS source code it would be possible to slice the radio signal by both atmospheric depth and energy of the emitting particles simultaneously while also computing histograms of their kinematics in order to study this hypothesis directly. In the more distant future it may be possible to gradually transform our templates closer and closer to unity or potentially eliminate them entirely turning our synthesis model into a fully analytical one. As some measure of shower-to-shower fluctuation is very real due to quantum particle physics a statistical extension model would still be necessary to fully match CoREAS even in this case.

The demonstrated potential of a fast template-based synthesis model and avenues for future improvements should also serve to motivate investment in next-generation dense and precise radio air shower experiments such as the efforts to utilise the SKA-low radio astronomy interferometer [6]. Synergies are readily apparent: our model can tackle the vast number of antenna positions without exorbitant demands on computation time even for high-precision sampling methods [5] or interferometric inference. On the other hand direct hits on such an array would provide excellent tests of our model as well as our theoretical understanding of radio emission from cosmic ray air showers.

BIBLIOGRAPHY

[1] M. G. Aartsen et al. "Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector". In: *Science* 342 (2013), p. 1242856.

[2] B. P. Abbott et al. "Observation of Gravitational Waves from a Binary Black Hole Merger". In: *Phys. Rev. Lett.* 116.6 (2016), p. 061102.

[3] M. P. van Haarlem et al. "LOFAR: The LOw-Frequency ARray". In: 556, A2 (Aug. 2013), A2. DOI: `10.1051/0004-6361/201220873`. arXiv: `1305.3550 [astro-ph.IM]`.

[4] P. Schellart et al. "Detecting cosmic rays with the LOFAR radio telescope". In: *Astron. Astrophys.* 560 (2013), A98.

[5] S. Buitink et al. "Method for high precision reconstruction of air shower X_{max} using two-dimensional radio intensity profiles". In: *Phys. Rev.* D90.8 (2014), p. 082003.

[6] Tim Huege et al. "Ultimate precision in cosmic-ray radio detection — the SKA". In: *EPJ Web Conf.* 135 (2017), p. 02003.

[7] Torsten Enßlin. "Information field theory". In: *AIP Conf. Proc.* 1553 (2013), p. 184.

[8] A. M. Hillas. "The Origin of Ultra-High-Energy Cosmic Rays". In: *Ann. Rev. Astron. Astrophys.* 22 (Jan. 1984), pp. 425–444.

[9] A. R. Bell. "Cosmic ray acceleration". In: *Astroparticle Physics* 43 (Mar. 2013), pp. 56–70.

[10] A. V. Olinto, K. Kotera, and D. Allard. "Ultrahigh Energy Cosmic Rays and Neutrinos". In: *Nucl. Phys. Proc. Suppl.* 217 (2011), pp. 231–236.

[11] V. S. Beresinsky and G. T. Zatsepin. "Cosmic rays at ultra high energies (neutrino?)" In: *Physics Letters B* 28.6 (Jan. 1969), pp. 423–424.

[12] M. G. Aartsen et al. "Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A". In: *Science* 361.6398 (2018), eaat1378.

[13] Frank G. Schröder. "News from Cosmic Ray Air Showers (ICRC 2019 – Cosmic Ray Indirect Rapport)". In: *PoS* ICRC2019 (2019), p. 030.

[14] S. P. Swordy. "The Energy Spectra and Anisotropies of Cosmic Rays". In: *Space Science Reviews* 99 (Oct. 2001), pp. 85–94.

[15] Bruno Rossi and Kenneth Greisen. "Cosmic-Ray Theory". In: *Reviews of Modern Physics* 13.4 (Oct. 1941), pp. 240–309.

[16] Karl-Heinz Kampert and Michael Unger. "Measurements of the cosmic ray composition with air shower experiments". In: *Astroparticle Physics* 35.10 (May 2012), pp. 660–678.

[17] KASCADE Collaboration et al. "Muon density measurements with the KASCADE central detector". In: *Astroparticle Physics* 16.4 (Feb. 2002), pp. 373–386.

[18] Markus Ahlers and Francis Halzen. "Opening a new window onto the universe with IceCube". In: *Progress in Particle and Nuclear Physics* 102 (Sept. 2018), pp. 73–88.

[19] D. Heck et al. "CORSIKA: A Monte Carlo code to simulate extensive air showers". In: *FZKA-6019* (1998).

[20] T. K. Gaisser and A. M. Hillas. "Reliability of the Method of Constant Intensity Cuts for Reconstructing the Average Development of Vertical Showers". In: *International Cosmic Ray Conference*. Vol. 8. International Cosmic Ray Conference. Jan. 1977, p. 353.

[21] S. J. Sciutto. "AIRES: A system for air shower simulations". In: *arXiv e-prints*, astro-ph/9911331 (Nov. 1999), astro–ph/9911331.

[22] T. Huege. "Radio detection of cosmic ray air showers in the digital era". In: *Phys. Rept.* 620 (2016), pp. 1–52.

[23] Anne Zilles. "Modeling of radio emission from a particle cascade in a magnetic field and its experimental validation". In: *EPJ Web Conf.* 135 (2017), p. 01018.

[24] Christian Glaser et al. "Simulation of radiation energy release in air showers". In: *Journal of Cosmology and Astroparticle Physics* 2016.9, 024 (Sept. 2016), p. 024.

[25] T Antoni et al. "The Cosmic ray experiment KASCADE". In: *Nucl. Instrum. Meth.* A513 (2003), pp. 490–510.

[26] S. Ostapchenko. "Monte Carlo treatment of hadronic interactions in enhanced Pomeron scheme: QGSJET-II model". In: *Phys. Rev. D* 83 (1 Jan. 2011), p. 014018.

[27] Giuseppe Battistoni et al. "The FLUKA Code: An Accurate Simulation Tool for Particle Therapy". In: *Frontiers in Oncology* 6 (2016), p. 116.

[28] *U.S. standard atmosphere, 1976*. Tech. rep. Oct. 1976.

[29] T. Bergmann et al. "One-dimensional hybrid approach to extensive air shower simulation". In: *Astroparticle Physics* 26.6 (Jan. 2007), pp. 420–432.

[30] The Pierre Auger Collaboration. "The Pierre Auger Cosmic Ray Observatory". In: *arXiv e-prints*, arXiv:1502.01323 (Feb. 2015).

[31] R. A. Ong. "Very high-energy gamma-ray astronomy". In: *Physics Reports* 305.3-4 (Jan. 1998), pp. 93–202.

[32] The Pierre Auger Collaboration. "The Fluorescence Detector of the Pierre Auger Observatory". In: *arXiv e-prints*, arXiv:0907.4282 (July 2009).

[33] A. G. Vieregg, K. Bechtol, and A. Romero-Wolf. "A technique for detection of PeV neutrinos using a phased radio array". In: *JCAP* 2016.2, 005 (Feb. 2016), p. 005. DOI: `10.1088/1475-7516/2016/02/005`.

[34] Pierre Auger Collaboration. "The exposure of the hybrid detector of the Pierre Auger Observatory". In: *Astroparticle Physics* 34.6 (Jan. 2011), pp. 368–381.

[35] G. A. Askar'yan. "Excess negative charge of an electron-photon shower and its coherent radio emission". In: *Sov. Phys. JETP* 14.2 (1962). [Zh. Eksp. Teor. Fiz.41,616(1961)], pp. 441–443.

[36] O. Scholten, K. Werner, and F. Rusydi. "A macroscopic description of coherent geo-magnetic radiation from cosmic-ray air showers". In: *Astroparticle Physics* 29.2 (Mar. 2008), pp. 94–103.

[37] H. Schoorlemmer. "Tuning in on cosmic rays". Ph. D. Radboud Universiteit Nijmegen, 2012.

[38] T. Huege, M. Ludwig, and C. W. James. "Simulating radio emission from air showers with CoREAS". In: *AIP Conf. Proc.* 1535 (2013), p. 128.

[39] O. Scholten et al. "Measurement of the circular polarization in radio emission from extensive air showers confirms emission mechanisms". In: *Phys. Rev. D* 94.10, 103010 (Nov. 2016), p. 103010.

[40] P. Schellart et al. "Polarized radio emission from extensive air showers measured with LOFAR". In: *Journal of Cosmology and Astroparticle Physics* 2014.10, 014 (Oct. 2014), p. 014.

[41] O. Scholten et al. "Analytic calculation of radio emission from parametrized extensive air showers: A tool to extract shower parameters". In: *Phys. Rev. D* 97.2, 023005 (Jan. 2018), p. 023005.

[42] T. Huege and H. Falcke. "Radio emission from cosmic ray air showers. Monte Carlo simulations". In: *Astronomy and Astrophysics* 430 (Feb. 2005), pp. 779–798.

[43] Marvin Gottowik et al. "Determination of the absolute energy scale of extensive air showers via radio emission: Systematic uncertainty of underlying first-principle calculations". In: *Astroparticle Physics* 103 (Dec. 2018), pp. 87–93.

[44] Jaime Alvarez-Muniz, Washington R. Carvalho Jr., and Enrique Zas. "Monte Carlo simulations of radio pulses in atmospheric showers using ZHAireS". In: *Astropart. Phys.* 35 (2012), pp. 325–341.

[45] Clancy W. James et al. "General description of electromagnetic radiation processes based on instantaneous charge acceleration in 'endpoints'". In: *Phys. Rev.* A84 (2011), p. 056602.

[46] Frank Nerling et al. "Universality of electron distributions in high-energy air showers: Description of Cherenkov light production". In: *Astropart. Phys.* 24 (2006), pp. 421–437.

[47] S. Lafebre et al. "Universality of electron-positron distributions in extensive air showers". In: *Astropart. Phys.* 31 (2009), pp. 243–254.

[48] The Pierre Auger Collaboration. "Data-driven estimation of the invisible energy of cosmic ray showers with the Pierre Auger Observatory". In: *Phys. Rev. D* 100.8, 082003 (Oct. 2019), p. 082003.

[49] Frank G. Schröder. "Radio detection of cosmic-ray air showers and high-energy neutrinos". In: *Progress in Particle and Nuclear Physics* 93 (Mar. 2017), pp. 1–68.

[50] M. Gottowik. "Measurements of Inclined Air Showers with the Auger Engineering Radio Array at the Pierre Auger Observatory". In: *36th International Cosmic Ray Conference (ICRC2019)*. Vol. 36. International Cosmic Ray Conference. July 2019.

[51] T. Huege, F. Schlüter, and L. Brenk. "Symmetrizing the signal distribution of radio emission from inclined air showers". In: *36th International Cosmic Ray Conference (ICRC2019)*. Vol. 36. International Cosmic Ray Conference. July 2019, p. 294.

[52] E. Holt. "Simulationsstudie für ein großskaliges Antennenfeld zur Detektion von Radioemission ausgedehnter Luftschauer". Diploma Thesis. KIT, 2013.

12

ACKNOWLEDGEMENTS

I am deeply grateful to my acting supervisor Tim Huege for his constant insights and support despite not filling a formal position as a referee. No less gratitude goes to my referees Prof. Ralph Engel and P.D. Torsten Enßlin for their specialised assistance when Tim and I ran out of ideas.

Special thanks to Prof. Olaf Scholten who welcomed me in Groningen and showed me the proper operation of MGMR3D in addition to providing constant advice based on his years of experience in designing analytical radio models.

Honourable mention to Prof. Guido Drexlin for acting as my initial referee before Prof. Engel became available and to Prof. Anke-Susanne Müller whom I thankfully had no need to fall back on in her position as my mentor.

I would also like to thank Dr. Andreas Haungs as the leader of the CR Technologies group for welcoming me as well as (now) Prof. Frank Schröder and all other group members for productive discussions not to mention a friendly and welcoming atmosphere at our institute.

The support of Dr. Tanguy Pierog and Dr. Dieter Heck has been invaluable in operating, modifying and fixing bugs within the CORSIKA code framework as my simulation needs were by no means standard fare. Thanks again to Tim for providing the same support for CoREAS specifically.

But no productive work would be possible without the immediate relief brought by office mates and corridor kin! Qader, my first day here at KIT was your last yet your office plant is still going strong. Olga, you showed me the fundamentals of being a Ph. D. student and also what to expect in the final stages. Ewa, you were great company and I miss your cheerful nature and colourful workplace decorations. Felix, I hope you are doing well in Argentina and return in time for a proper farewell. Roxanne, you only just got started and are already fighting to bring CR-radio to one of the most remote places on Earth.

Aswathi, Agnieszka and Marie: you girls were awesome! Vladimir, Victoria, Hrvoje, Paras: you too! Max, you narrowly beat me to the finish line. Tom, you or Agnieszka are next! Donghwa, Katrin, I hope you are not too lonely *aaall* the way over in the other building. Good luck and have fun all of you! Dima, Louis, there were so many great people here over the years I cannot list them all.

@Doris: at least I have not crashed the grid. yet. You, Sabine and the other supporting staff do incredible work making everything run smoothly so we can focus on the science.