



Durham E-Theses

Precision Z Boson Phenomenology at the LHC

MORGAN, THOMAS,ANTHONY

How to cite:

MORGAN, THOMAS,ANTHONY (2016) *Precision Z Boson Phenomenology at the LHC*, Durham theses, Durham University. Available at Durham E-Theses Online: <http://etheses.dur.ac.uk/11773/>

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

Precision Z Boson Phenomenology at the LHC

Thomas Anthony Morgan

A Thesis presented for the degree of
Doctor of Philosophy



Institute for Particle Physics Phenomenology
Department of Physics
Durham University
United Kingdom

September 2016

Precision Z Boson Phenomenology at the LHC

Thomas Anthony Morgan

Submitted for the degree of Doctor of Philosophy

September 2016

Abstract: In this thesis the Next-to-Next-to-Leading Order (NNLO) perturbative QCD corrections to Z boson production in association with hard QCD radiation are considered, both where the hard QCD radiation forms an observable jet and where no jet definition is required. The Infrared (IR) divergences are regulated using the antenna subtraction method and all possible initial state configurations and colour contributions are considered. This research is the first complete NNLO calculation relevant for LHC phenomenology using the antenna subtraction method. The work in this thesis forms the backbone of the NNLOJET framework, a highly optimised and flexible Monte Carlo integrator constructed to perform precision LHC phenomenology. The results of this computation are presented both as total cross sections and differential distributions for a wide range of LHC observables. Excellent agreement is observed between the experimental data from ATLAS and CMS for the fully inclusive normalised transverse momentum distribution and the NNLO prediction, opening the possibility of using this calculation to constrain the gluon parton distribution function.

Declaration

The work in this thesis is based on research carried out in the Institute for Particle Physics Phenomenology, Department of Physics, Durham University, United Kingdom. No part of this thesis has been submitted elsewhere for any other degree or qualification. All work is my own in collaboration with my supervisor Professor E.W.N. Glover, unless referenced to otherwise within the text. Chapters 3 and 6 are based upon research done in collaboration with A. Gehrmann-De Ridder, T. Gehrmann and A. Huss in the publications,

A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss, and T. A. Morgan. “Precise QCD predictions for the production of a Z boson in association with a hadronic jet”. *Phys. Rev. Lett.* 117.2 (2016), p. 022001. arXiv: 1507.02850 [hep-ph],

A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss, and T. A. Morgan. “The NNLO QCD corrections to Z boson production at large transverse momentum”. *JHEP* 07 (2016), p. 133. arXiv: 1605.04295 [hep-ph],

and research presented as a Proceedings in Science,

A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss, and T. A. Morgan. “Z+jet production at NNLO”. *13th DESY Workshop on Elementary Particle Physics: Loops and Legs in Quantum Field Theory (LL2016) Leipzig, Germany, April 24-29, 2016*. 2016. arXiv: 1607.01749 [hep-ph].

Chapter 5 is based upon research in collaboration with A. Gehrmann-De Ridder, T. Gehrmann and A. Huss presented as Proceedings in Science,

A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss, and T. A. Mor-

gan. “NNLO QCD corrections for Z boson plus jet production”. *Proceedings, 12th International Symposium on Radiative Corrections (Radcor 2015) and LoopFest XIV (Radiative Corrections for the LHC and Future Colliders)*. 2016. arXiv: 1601.04569 [hep-ph]

Copyright © September 2016 by Thomas Anthony Morgan.

“The copyright of this thesis rests with the author. No quotations from it should be published without the author’s prior written consent and information derived from it should be acknowledged.”

Acknowledgements

Firstly, I would like to thank my supervisor Nigel Glover for his incredible insight and seemingly infinite patience over the last four years. His guidance throughout this project helped us overcome countless problems I believed were insurmountable. For that I will be forever grateful.

A special thanks must be reserved for my collaborator Alexander Huss. The vast number of Skype exchanges and his many contributions to the NNLOJET framework have been indispensable over the last few years.

Of course none of this would be possible without the hard work and support provided by everyone within the NNLOJET collaboration: Xuan Chen, Juan Cruz-Martinez, James Currie, Aude Gehrmann-De Ridder, Thomas Gehrmann, Jan Niehues, João Pires, Duncan Walker and Steven Wells. They have all contributed so much to my four years in this collaboration with both interesting discussions and wonderful insights of the field of QCD.

My sincere gratitude goes out to Oliver Smith and Jeppe Anderson, without whom I could not have possibly completed my campaign of consuming every possible CPU hour available to me. Also thank you to Daniela Bauer and Simon Fayer who have provided a huge amount of support for the Dirac submission system. It was always a pleasure being the ‘pheno canary’.

Thank you to all of the PhD students, both past and present, who have made the last four years a joy. Life just won’t be the same without Gilberto’s crazy stories of life in Mexico, Age of Empires sessions and spontaneous visits to the pub.

I would never have made it this far if it were not for my friends and family for keeping me sane all these years. Thank you to the best housemates anyone could ask for, Matthew Jones and Ben Dobson, for providing 3 years filled with Battlestar, Starcraft and everything else nerdy. Tricity misses you both almost as much as I do.

Finally, I would like to thank my partner Dean. Nothing fills me with as much joy as knowing that these final few words of this chapter of my life mark the beginning of our next chapter together.

This research was supported by the UK Science and Technology Facilities Council and in part by the Research Executive Agency (REA) of the European Union through the ERC Advanced Grant MC@NNLO (340983).

Contents

| | |
|---|------------|
| Abstract | ii |
| List of Figures | ix |
| List of Tables | xix |
| 1 Introduction to Standard Model physics | 3 |
| 1.1 Lagrangian dynamics and gauge symmetries | 4 |
| 1.2 Electroweak interactions | 8 |
| 1.3 The Higgs mechanism | 11 |
| 1.4 Renormalisation of QCD | 17 |
| 1.5 Colour decomposition of matrix elements | 23 |
| 1.6 Infrared singularities | 25 |
| 1.6.1 Virtual IR singularities at one-loop | 25 |
| 1.6.2 Virtual IR singularities at two-loops | 27 |
| 1.6.3 Real IR singularities at tree level | 29 |
| 1.6.4 Real IR singularities at one-loop | 33 |
| 1.6.5 Double real IR singularities at tree level | 33 |
| 1.7 Hadron-hadron collisions and cross sections | 36 |
| 1.7.1 The improved parton model | 39 |
| 1.7.2 Regulating IR divergences in cross sections | 42 |
| 1.8 Jet algorithms and IR safe observables | 46 |
| 1.8.1 Cone algorithms | 48 |
| 1.8.2 Sequential recombination algorithms | 50 |

| | | |
|----------|---|-----------|
| 2 | Antenna Subtraction | 52 |
| 2.1 | Unintegrated antenna functions | 52 |
| 2.1.1 | Sub-antennae | 60 |
| 2.2 | Integrated antenna functions | 62 |
| 2.3 | Antenna subtraction at NLO | 64 |
| 2.3.1 | Real subtraction | 64 |
| 2.3.2 | Virtual subtraction | 65 |
| 2.4 | Antenna subtraction at NNLO | 69 |
| 2.4.1 | Double real subtraction | 69 |
| 2.4.2 | Real virtual subtraction | 75 |
| 2.4.3 | Double virtual subtraction | 78 |
| 2.5 | D_4^0 antenna | 82 |
| 2.5.1 | Gluon initiated D_4^0 antenna | 83 |
| 2.6 | Antenna subtraction with processes involving a vector boson | 86 |
| 3 | Drell–Yan phenomenology at the LHC | 90 |
| 3.1 | Motivation for precision predictions of Drell–Yan observables | 91 |
| 3.2 | Theoretical corrections to Drell–Yan production | 93 |
| 3.3 | Drell–Yan production in association with hadronic jets | 95 |
| 3.4 | Transverse momentum distribution of the Z boson | 96 |
| 4 | Numerical implementation of NNLO calculations | 99 |
| 4.1 | Numerical testing of subtraction terms | 99 |
| 4.2 | Virtual numerical integration | 103 |
| 4.2.1 | Plus distributions | 105 |
| 4.3 | Scale dependence and the running of $\alpha_s(\mu_R^2)$ | 110 |
| 4.4 | Double real numerical integration | 113 |
| 4.4.1 | Double real phase space regions | 115 |
| 4.4.2 | Double real phase space sampling | 119 |
| 4.5 | Technical cut dependence of real integrals | 123 |
| 4.5.1 | Reduced phase space technical cut | 124 |
| 4.6 | NNLO automation techniques | 128 |
| 4.6.1 | Auto-generation of subtraction terms | 129 |

| | | |
|----------|---|------------|
| 4.6.2 | Auto-generation of driver routines | 130 |
| 4.6.3 | Channel breakdown of Z +jet at NNLO | 133 |
| 4.7 | Numerical optimisation | 137 |
| 4.7.1 | Storage element | 138 |
| 4.7.2 | Colour sampling | 141 |
| 4.7.3 | Reweighting | 143 |
| 4.7.4 | Multiscale integration | 144 |
| 4.7.5 | Observable matching | 145 |
| 5 | Analysing data from NNLO calculations | 147 |
| 5.1 | Statistical Outliers | 147 |
| 5.2 | Weighted averaging | 148 |
| 6 | Phenomenological results | 153 |
| 6.1 | NNLO corrections to Z +jet production at the LHC | 153 |
| 6.2 | Transverse momentum distribution of the Z boson | 159 |
| 6.2.1 | Double-differential distributions | 162 |
| 7 | Conclusions | 171 |
| | Bibliography | 175 |
| A | X_3^0 Antenna functions | 191 |
| A.1 | $\gamma^* \rightarrow q\bar{q}$ antenna functions | 191 |
| A.2 | $\tilde{\chi} \rightarrow \tilde{g}g$ antenna functions | 192 |
| A.3 | $H \rightarrow gg$ antenna functions | 193 |
| B | Validation plots and distributions | 194 |
| B.1 | NLO validation | 194 |
| B.1.1 | Z +jet at NLO | 194 |
| B.1.2 | Z +2 jets at NLO | 197 |
| B.2 | Technical cut plots | 202 |

List of Figures

| | | |
|-----|--|----|
| 1.1 | The ‘Mexican hat’ Higgs potential for a single complex scalar ϕ , where $V(\phi) = \mu^2\phi^*\phi - \lambda(\phi^*\phi)^2$ and $\mu^2, \lambda < 0$. The red line denotes the set of field configurations for ϕ which minimises the potential. | 13 |
| 1.2 | A diagrammatic representation of the colour flow in the Fierz identity. In the first term on the right hand side there is a colour exchange between the quark lines. In the second term the gluon does not exchange colour. | 24 |
| 1.3 | A Feynman diagram of a gluon emission from a fermion line. The blob denotes the rest of the hard scattering process. | 29 |
| 1.4 | Colour connected double unresolved emissions for $q\bar{q} \rightarrow gg$. The vertical dashed lines denote unresolved gluon emissions. | 34 |
| 1.5 | Almost colour connected double unresolved emissions for $q\bar{q} \rightarrow ggg$. The vertical dashed lines denote unresolved gluon emissions. | 35 |
| 1.6 | Colour disconnected double unresolved emissions for $q\bar{q} \rightarrow ng$. The vertical dashed lines denote unresolved gluon emissions. | 36 |
| 1.7 | An example PDF set from NNPDF2.3 [64] for a proton with $\mu_F = 10$ GeV. The gluon PDF has been reduced by a factor of 10 for clarity. | 37 |
| 1.8 | Emission of a parton on an initial state leg, changing the instantaneous momentum fraction of the initial state parton. | 39 |
| 1.9 | A schematic of collinear splitting with seeded cone algorithms that use seeds above a partonic kinematic threshold. A hard parton (red) can undergo collinear splitting into two partons which are both below the kinematic threshold, splitting a larger jet into two smaller jets. | 48 |

| | | |
|------|---|-----|
| 1.10 | A schematic of soft emissions with seeded cone algorithms that use all parton momenta as seeds. A soft emission (red) between two hard partons merges the two jets into one larger jet. | 49 |
| 2.1 | A schematic of the contributing terms to the double real (red), real virtual (blue) and double virtual (green) subtraction terms using the antenna formalism. The arrows indicate where terms that are introduced are subsequently integrated out. | 81 |
| 3.1 | The leading order contribution to the Drell–Yan cross section in the naïve parton model. X denotes the hadronic remnants of the event. . . | 91 |
| 3.2 | Feynman diagrams depicting (a) a real correction to Drell–Yan production from a qg initiated channel at $\mathcal{O}(\alpha_s)$ and (b) a double real correction to Drell–Yan production from a gg initiated channel at $\mathcal{O}(\alpha_s^2)$ | 94 |
| 3.3 | Representative Feynman diagrams for (a) two-loop Z boson plus three parton amplitudes (b) one-loop Z boson plus four parton amplitudes and (c) tree level Z boson plus five parton amplitudes. | 95 |
| 3.4 | A schematic diagram demonstrating the Z boson recoiling against hard radiation. | 97 |
| 4.1 | Example spike plots for the triple gluon collinear limit leading colour contribution to the $q\bar{q} \rightarrow ggg$ for Z +jet at NNLO. x_i is a small dimensionless parameter that controls the proximity of the phase space to the triple collinear limit. (a) shows the limit without an azimuthal rotation and (b) shows the limit with an azimuthal rotation. The x axis differs between the two plots for clarity. | 102 |
| 4.2 | The $\sigma'(\mu_R, \alpha_s(\mu_R^2))$ cross section at NNLO accuracy for Z +jet. The green band denotes the prediction from Eq. (4.3.6) running from $\mu_0 = M_Z$, where the area of the band is the statistical uncertainty. The red data points are taken directly from the NNLOJET code ran with $\mu_F = M_Z$ and $\mu_R = [1/2, 1, 2] \cdot M_Z$, with the statistical uncertainty shown. . . | 113 |
| 4.3 | The rotation by $\pi/2$ applied in the phase space generator around the direction of the initial state parton p'_1 in the rest frame of p_{45} | 117 |

- 4.4 Spike plots for a triple collinear initial-final-final $\hat{1}_g||3_g||4_g$ limit for the $qg \rightarrow qggZ$ leading colour contribution to Z +jet production with (a) no rotations applied, (b) a rotation about the p_1 axis and (c) a rotation about the shared $p_1 + p_3 + p_4$ collinear axis, boosted back on beam. The phase space points generated before the rotation is applied are identical in the three spike plots. 118
- 4.5 Contributions to the cross section for different values of the full technical cut, y_0 , for (a) double real region a and (b) double real region b , where the error bars on the datapoints are statistical. Approximately the same number of events were used for each datapoint ($\mathcal{O}(10^8)$). The green line represents an equivalent calculation with significantly more events ($\mathcal{O}(10^9)$) for $y_0 = 10^{-6}$, where the green band denotes its corresponding statistical uncertainty. 127
- 4.6 Contributions to the cross section for different values of δ_0 for (a) double real region a and (b) double real region b , where the error bars on the datapoints are statistical. Approximately the same number of events were used for each datapoint ($\mathcal{O}(10^8)$) and the full technical cut was set to $y_0 = 10^{-6}$. The green line represents an equivalent calculation with significantly more events ($\mathcal{O}(10^9)$) for $y_0 = 10^{-6}$ and $\delta_0 = 0$, where the green band denotes its corresponding statistical uncertainty. 127
- 4.7 A flowchart of the processing required to obtain numerical results using NNLOJET. The red components represent a time consuming process in terms of the implementation. The orange components indicate the numerically time consuming contributions. The dashed arrows represent a feedback loop where the construction of the ‘map’ file is determined from the quality of the spike plots and/or the cancellation of ϵ poles. 133
- 5.1 The total cross section $\sigma(N/k, k \cdot M)$ at NNLO as a function of k , the number of runs combined to form a pseudorun with $k \cdot M$ events. The cross section is then the weighted average of the pseudoruns. The errors on the individual data points are statistical. 150

| | | |
|-----|---|-----|
| 5.2 | The integral of the y^Z distribution, $\sigma(N/k, k \cdot M)$ at NNLO as a function of k . The errors on the individual data points are statistical. The green band denotes the total cross section evaluated at $k = 100$ with its statistical uncertainty. | 151 |
| 5.3 | The ratio of the NNLO against the NLO prediction for individual bins of the rapidity distribution of the Z boson for different values of k . The grey bands are the unweighted result. The errors on the individual data points are statistical. | 151 |
| 5.4 | The ratio of the NNLO against the NLO prediction for individual bins of the p_T distribution of the Z boson for different values of k . The grey bands are the unweighted result. The errors on the individual data points are statistical. | 152 |
| 6.1 | The transverse momentum distribution of the leading jet in inclusive Z +jet production in pp collisions with $\sqrt{s} = 8$ TeV at LO (blue), NLO (green), NNLO (red). The lower panel shows the perturbative K factors for NLO/LO (turquoise) and NNLO/NLO (mauve). | 155 |
| 6.2 | The rapidity distribution of the leading jet in inclusive Z +jet production in pp collisions with $\sqrt{s} = 8$ TeV at LO (blue), NLO (green), NNLO (red). The lower panel shows the perturbative K factors for NLO/LO (turquoise) and NNLO/NLO (mauve). | 156 |
| 6.3 | The transverse momentum distribution of the Z boson in inclusive Z +jet production in pp collisions with $\sqrt{s} = 8$ TeV at LO (blue), NLO (green), NNLO (red). The lower panel shows the perturbative K factors for NLO/LO (turquoise) and NNLO/NLO (mauve). For the NLO and NNLO predictions, the two bins extending $p_T^Z \in [20, 40]$ GeV have been merged into one. | 156 |
| 6.4 | The rapidity distribution of the Z boson in inclusive Z +jet production in pp collisions with $\sqrt{s} = 8$ TeV at LO (blue), NLO (green), NNLO (red). The lower panel shows the perturbative K factors for NLO/LO (turquoise) and NNLO/NLO (mauve). | 157 |

-
- 6.5 The transverse momentum distribution for the negatively charged lepton for Z +jet production in pp collisions with $\sqrt{s} = 8$ TeV at LO (blue), NLO (green), NNLO (red). The lower panel shows the perturbative K factors for NLO/LO (turquoise) and NNLO/NLO (mauve). 158
- 6.6 The pseudorapidity distribution for the negatively charged lepton for Z +jet production in pp collisions with $\sqrt{s} = 8$ TeV at LO (blue), NLO (green), NNLO (red). The lower panel shows the perturbative K factors for NLO/LO (turquoise) and NNLO/NLO (mauve). 158
- 6.7 The unnormalised Z boson transverse momentum distribution for the cuts given in Tab. 6.2 and $66 \text{ GeV} < m_{\ell\ell} < 116 \text{ GeV}$. ATLAS data is taken from Ref. [143]. The luminosity error is not shown. The green bands denote the NLO prediction with scale uncertainty and the blue bands show the NNLO prediction with scale uncertainty. 161
- 6.8 The normalised Z boson transverse momentum distribution for the cuts given in Tab. 6.2 and $66 \text{ GeV} < m_{\ell\ell} < 116 \text{ GeV}$. ATLAS data is taken from Ref. [143]. The green bands denote the NLO prediction with scale uncertainty and the blue bands show the NNLO prediction with scale uncertainty. 162
- 6.9 The unnormalised double-differential transverse momentum distribution for the Z boson in windows of invariant mass of the leptons, $m_{\ell\ell}$, with a rapidity cut on the Z boson of $|y^Z| < 2.4$. The ATLAS data is taken from Ref. [143]. The luminosity error is not shown. The green bands denote the NLO prediction with scale uncertainty and the blue bands show the NNLO prediction with scale uncertainty. 163

- 6.10 The inclusive dilepton cross section for the same $m_{\ell\ell}$ bins as in Fig. 6.9 and with a rapidity cut on the Z boson of $|y^Z| < 2.4$. The experimental data is taken from the ATLAS analysis in Ref. [143]. The ticks on the vertical error bands denote the systematic uncertainty from the measurement, the vertical bars without the ticks are the luminosity uncertainty only. The blue bands show the NNLO prediction with scale uncertainty. The theoretical cross section predictions for the invariant mass windows $12 < m_{\ell\ell} < 20$ GeV, $20 < m_{\ell\ell} < 30$ GeV, and $30 < m_{\ell\ell} < 46$ GeV are calculated using the Z +jet calculation since the additional kinematical cut $p_T^Z > 45$ GeV prevents contributions from the LO Drell–Yan configuration in these invariant mass windows. 164
- 6.11 The normalised double-differential transverse momentum distribution for the Z boson in windows of invariant mass of the leptons, $m_{\ell\ell}$, with a rapidity cut on the Z boson of $|y^Z| < 2.4$. The ATLAS data is taken from Ref. [143]. The green bands denote the NLO prediction with scale uncertainty and the blue bands show the NNLO prediction with scale uncertainty. The normalisation cross section for the invariant mass windows $12 < m_{\ell\ell} < 20$ GeV, $20 < m_{\ell\ell} < 30$ GeV, and $30 < m_{\ell\ell} < 46$ GeV are calculated using the Z +jet calculation since the kinematical cut $p_T^Z > 45$ GeV prevents contributions from the LO Drell–Yan configuration in these invariant mass windows. 165
- 6.12 The unnormalised double-differential transverse momentum distribution for the Z boson in windows of rapidity of the Z boson, y^Z , with an invariant mass cut on the final state leptons of $66 \text{ GeV} < m_{\ell\ell} < 116 \text{ GeV}$. The ATLAS data is taken from Ref. [143]. The luminosity error is not shown. The green bands denote the NLO prediction with scale uncertainty and the blue bands show the NNLO prediction. 166

- 6.13 The inclusive dilepton cross section for the same $|y^Z|$ bins as in Fig. 6.12 and an invariant mass cut on final state leptons of $66 \text{ GeV} < m_{\ell\ell} < 116 \text{ GeV}$. The ATLAS data is extracted from Ref. [143] by summing up the transverse momentum distributions in the respective bins. The vertical error bars are given by the luminosity uncertainty. The blue bands show the NNLO prediction with scale uncertainty. 167
- 6.14 The normalised double-differential transverse momentum distribution for the Z boson in windows of rapidity of the Z boson, y^Z , with an invariant mass cut on final state leptons of $66 \text{ GeV} < m_{\ell\ell} < 116 \text{ GeV}$. The ATLAS data is taken from Ref. [143]. The green bands denote the NLO prediction with scale uncertainty and the blue bands show the NNLO prediction. 168
- 6.15 The unnormalised double-differential transverse momentum distribution for the Z boson in windows of rapidity of the Z boson, y^Z , with an invariant mass cut on final state leptons of $81 \text{ GeV} < m_{\ell\ell} < 101 \text{ GeV}$. The CMS data is taken from Ref. [145]. The experimental error bars include the luminosity uncertainty. The green bands denote the NLO prediction with scale uncertainty and the blue bands show the NNLO prediction. 169
- 6.16 The normalised double-differential transverse momentum distribution for the Z boson in windows of rapidity of the Z boson, y^Z , with an invariant mass cut on final state leptons of $81 \text{ GeV} < m_{\ell\ell} < 101 \text{ GeV}$. The CMS data is taken from Ref. [145]. The green bands denote the NLO prediction with scale uncertainty and the blue bands show the NNLO prediction. 170
- B.1 The transverse momentum distribution of the leading jet in inclusive Z +jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical. 195

-
- B.2 The rapidity distribution of the leading jet in inclusive Z +jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical. 196
- B.3 The transverse momentum distribution of the Z boson in inclusive Z +jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical. 197
- B.4 The rapidity distribution of the Z boson in inclusive Z +jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical. 198
- B.5 The transverse momentum distribution of the negatively charged lepton in inclusive Z +jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical. 198
- B.6 The rapidity distribution of the negatively charged lepton in inclusive Z +jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical. 199
- B.7 The transverse momentum distribution of the leading jet in inclusive Z +2 jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical. 200

- B.8 The rapidity distribution of the leading jet in inclusive $Z+2$ jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical. 201
- B.9 The transverse momentum distribution of the Z boson in inclusive $Z+2$ jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical. 201
- B.10 The rapidity distribution of the Z boson in inclusive $Z+2$ jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical. 202
- B.11 The transverse momentum distribution of the negatively charged lepton in inclusive $Z+2$ jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical. 203
- B.12 The rapidity distribution of the negatively charged lepton in inclusive $Z+2$ jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical. 203
- B.13 Contributions to the cross section for Z +jet production for different values of the technical cut, y_0 , for (a) real corrections at NLO and (b) real virtual corrections at NNLO. The error bars on the datapoints are statistical. 204

- B.14 The transverse momentum distribution for the leading jet, double real region a contribution normalised to the high statistics double real region a result used in section 6.1. The errors on each datapoint are statistical. The technical cut values are (a) $y_0 = 1 \cdot 10^{-5}$, (b) $y_0 = 3 \cdot 10^{-6}$, (c) $y_0 = 1 \cdot 10^{-6}$, (d), $y_0 = 3 \cdot 10^{-7}$, (e) $y_0 = 1 \cdot 10^{-7}$. The green band denotes the statistical uncertainty on the high statistics calculation. 205
- B.15 The transverse momentum distribution for the leading jet, double real region b contribution normalised to the high statistics double real region b result used in section 6.1. The errors on each datapoint are statistical. The technical cut values are (a) $y_0 = 1 \cdot 10^{-5}$, (b) $y_0 = 3 \cdot 10^{-6}$, (c) $y_0 = 1 \cdot 10^{-6}$, (d), $y_0 = 3 \cdot 10^{-7}$, (e) $y_0 = 1 \cdot 10^{-7}$. The green band denotes the statistical uncertainty from the high statistics calculation. 206
- B.16 The transverse momentum distribution for the leading jet, double real region a contribution normalised to the high statistics double real region a result used in section 6.1. The errors on each datapoint are statistical. The reduced technical cut values are (a) $\delta_0 = 10^{-3}$, (b) $\delta_0 = 10^{-4}$, (c) $\delta_0 = 10^{-5}$. The green band denotes the statistical uncertainty from the high statistics calculation. 207
- B.17 The transverse momentum distribution for the leading jet, double real region b contribution normalised to the high statistics double real region b result used in section 6.1. The errors on each datapoint are statistical. The reduced technical cut values are (a) $\delta_0 = 10^{-3}$, (b) $\delta_0 = 10^{-4}$, (c) $\delta_0 = 10^{-5}$. The green band denotes the statistical uncertainty from the high statistics calculation. 208

List of Tables

| | | |
|-----|--|----|
| 1.1 | The quantum numbers corresponding to the fermions in the GWS Lagrangian. T_w denotes the weak isospin, T_{w_3} denotes the component of T_w for W_μ^3 . Y is the hypercharge quantum number and Q is the resulting charge following spontaneous symmetry breaking. | 11 |
| 2.1 | A table of all of the tree level and one-loop antenna functions required for NNLO calculations [11]. Antenna denoted by a tilde are subleading colour contributions and antenna denoted with a hat are flavour dependent contributions. | 56 |
| 2.2 | A table of the desired and undesired limits of $D_3^0(1_q, 3_g, 4_g)$. It is important that in our antenna 1_q and 4_g act as the hard radiators and do not produce singularities as they go unresolved. | 60 |
| 2.3 | The definition of the final-final $\mathbf{J}_2^{(1)}$ dipoles in terms of integrated antenna and the corresponding singularity structure they subtract. | 66 |
| 2.4 | The definition of the initial-final $\mathbf{J}_2^{(1)}$ dipoles in terms of integrated antenna and the corresponding singularity structure they subtract. For brevity $\delta_2 = \delta(1 - x_2)$ | 67 |
| 2.5 | The definition of the initial-initial $\mathbf{J}_2^{(1)}$ dipoles in terms of integrated antenna and the corresponding singularity structure they subtract. For brevity $\delta_i = \delta(1 - x_i)$ for $i = 1, 2$ | 68 |
| 2.6 | The classification of double unresolved antenna without almost colour connected limits (X_4^0) and those with almost colour connected limits (\widetilde{X}_4^0) [68]. | 72 |
| 3.1 | The branching ratios for the dominant decay modes of the Z boson [25]. | 92 |

| | | |
|-----|---|-----|
| 4.1 | Approximate CPU hours each of the contributions to the full NNLO calculation of Z +jet. The double real phase space is divided into two regions, a and b , to optimise the sampling of distinct double unresolved configurations. | 114 |
| 4.2 | The number of combinations of reduced phase spaces for real corrections ($R(n)$) and double real corrections ($RR(n)$) to the process $pp \rightarrow n$ partons, where there are n partons in the leading order contribution. | 140 |
| 4.3 | A table of colour factors which are sampled for each initial state configuration separated by colour channel for the Z +jet double real integration. The importance of the sampling is in reverse numerical order (i.e. channel 1 is sampled most frequently and channel 5 is sampled the least often). In channels 3, 4 and 5 there are no longer any colour contributions for the gg and the $qQ + q\bar{Q} + \bar{q}\bar{Q}$ initial states, hence only the qg and gq channels are sampled for these channels. Contributions proportional to n_f are divided into n_{up} and n_{down} processes for Z +jet production. . . . | 142 |
| 6.1 | Channel breakdown of the total cross section for LO, NLO and NNLO for the scale choice $\mu = \mu_F = \mu_R = M_Z$. The theoretical uncertainty on each channel is estimated from the envelope of the $[1/2, 1, 2] \cdot M_Z$ scale choices. | 154 |
| 6.2 | Kinematical cuts used to define the fiducial phase space for the leading and subleading leptons, ordered in transverse momentum, for the measurements of ATLAS [143] and CMS [145]. | 160 |

Preface

The physics of particle colliders is a remarkable field of study. On the experimental side, the Large Hadron Collider (LHC) is a feat of modern engineering. Huge collaborations of experimental physicists and engineers are required to run, maintain and interpret the data. The discovery of a scalar particle at the ATLAS and CMS detectors in 2012 [5, 6] is truly a testament to a hugely successful experimental programme.

On the theoretical side, the unprecedented quality of the LHC data has led theorists to consider ever more complicated calculations to reduce theoretical uncertainties. The second iteration of the so called Les Houches ‘wishlist’ [7] demonstrates the colossal endeavour the theoretical community faces. Despite this, many of these calculations in the 2013 wishlist are either already complete or in progress.

The key question to ask is *why* is it important to reduce the uncertainties on theoretical calculations? It is clear that the Standard Model (SM) of particle physics is fundamentally flawed; it is unable to provide a suitable dark matter candidate or explain gravitational interactions. Nevertheless, it is clear that any deviations from the Standard Model observed by collider experiments will be subtle. The work done by theorists in providing precise Standard Model predictions is important in distinguishing Beyond the Standard Model (BSM) from Standard Model backgrounds.

The work in this thesis will focus on one important Standard Model background at the LHC, namely lepton pair production in association with a hadronic jet. This process is on the Les Houches wishlist and is important for distinguishing BSM signals such as supersymmetry. In particular, we will be including higher order perturbative corrections in the context of *Quantum Chromodynamics* (QCD) as an attempt to reduce theoretical uncertainties for this process. The terms calculated in this thesis are the full Next-to-Next-to-Leading Order (NNLO) corrections.

This calculation poses a significant challenge in terms of its analytical construction and numerical implementation. Indeed, it has been known for a long time that higher order terms in perturbative QCD contain high energy ultraviolet (UV) and low energy infrared (IR) divergences. UV divergences can be regulated through a renormalisation scheme choice, whereas IR divergences are far more subtle. The analysis presented by Bloch and Nordsieck [8] and the Kinoshita-Lee-Nauenberg (KLN) theorem [9, 10] demonstrate that all IR poles must cancel at each order in the perturbative series. However, how this cancellation arises is not immediately obvious and typically requires some form of subtraction to render the cancellation explicit. It is only after implementing a subtraction scheme that one is able to extract the physical, finite result.

The calculation presented in this work uses the *antenna subtraction* formalism [11–14], a flexible subtraction scheme which can be applied to a wide range of processes, including processes with hadronic initial states and final state jets. This particular process poses new challenges for the formalism however, including a high dimensional phase space and many subchannels that need to be integrated over. These will be discussed in detail within this thesis.

This thesis is structured as follows. Chapter 1 will form the introductory material to this research, focusing on basic concepts of Standard Model phenomenology. Chapter 2 will discuss the antenna subtraction formalism from the context of analytical expressions. Chapter 3 will then move onto a discussion of Z boson phenomenology and will motivate the necessity for precision for this particular LHC process. Chapters 4 and 5 will focus on the numerical aspects of the calculation. Finally, the numerical results of the calculation will be presented in chapter 6.

Chapter 1

Introduction to Standard Model physics

In this chapter we will review the basic concepts of Standard Model (SM) interactions within the context of a high energy collider environment, forming the basis for the research presented in this thesis. We begin with a review of $SU(N)$ non-abelian gauge theories from the context of Lagrangian dynamics. This will form the foundation of both QCD and electroweak (EW) sectors of the SM. This leads us naturally to construct the Glashow-Weinberg-Salam (GWS) [15–17] Lagrangian which describes EW interactions and explains the W and Z bosons mass terms using a Spontaneous Symmetry Breaking (SSB) mechanism.

For the strong sector of the Lagrangian, the approach of higher order corrections through a perturbative series of α_s will be introduced. The higher order corrections motivate us to consider renormalisation and mass factorisation techniques as a method to regulate singularities occurring naturally during computations from divergent loop diagrams and initial state collinear radiation respectively.

We will also introduce the concept of infrared (IR) singularity structures and approaches that can be applied to regulate these singularities. Given a suitable subtraction scheme to deal with the IR singularities, we have all the basic concepts required to begin calculating higher order QCD corrections to processes in a hadron-hadron collider environment.

1.1 Lagrangian dynamics and gauge symmetries

The Lagrangian density, henceforth referred to simply as the Lagrangian, is defined as the integrand of the space-time action,

$$S[\{\phi_i\}] = i \int d^4x \mathcal{L}(\{\phi_i(x)\}, \{\partial_\mu \phi_i(x)\}), \quad (1.1.1)$$

where $\{\phi_i\}$ denotes the set of fields with which the Lagrangian is defined. In the context of physical Quantum Field Theories (QFTs), the choice of Lagrangian must satisfy many basic constraints,

- the field definitions must be local, depending only on a single point in space-time,
- the Lagrangian must be Hermitian to ensure that the action is real,
- the Lagrangian must respect Poincaré invariance,
- all terms in the Lagrangian must be renormalisable.¹

The above constraints, whilst necessary, are far from sufficient to construct a theory capable of predicting physical observations. One must apply additional symmetries *a priori* that constrain the terms in the Lagrangian and the field content. In this section we will discuss the construction of a Lagrangian which is symmetric under local gauge transformations of an $SU(N)$ gauge group, where N is a constant yet to be defined. This discussion will underpin the foundation of QCD and EW interactions.

Firstly we introduce the spin- $\frac{1}{2}$ quark and anti-quark fields, $\psi(x)_{\alpha,i}$ and $\bar{\psi}_{\dot{\alpha},i}(x)$, which transform according to the fundamental and anti-fundamental representation of $SU(N)$ respectively,

$$\psi'_{\alpha,j}(x) \rightarrow U_{ij}(x)\psi_{\alpha,i}(x), \quad (1.1.2)$$

$$\bar{\psi}'_{\dot{\alpha},j}(x) \rightarrow \bar{\psi}_{\dot{\alpha},i}(x)U_{ij}^{-1}(x), \quad (1.1.3)$$

where $U(x)$ is the transformation operator belonging to the Lie algebra of $SU(N)$. The index i denotes the fundamental index of the gauge group and the index α denotes the

¹This will be discussed in further detail in section 1.4 but for the sake of this discussion leads to us only considering terms where the coupling has positive mass dimension.

spinor index of the field. Both of these sets of indices will be suppressed in further discussions.

In order to simultaneously provide a kinetic term for the quark fields and also satisfy the assumption that the Lagrangian is gauge invariant, it is also necessary to define a covariant derivative \not{D} ,

$$\not{D} \equiv \gamma^\mu (\partial_\mu - igA_\mu^a t^a), \quad (1.1.4)$$

where μ is the Lorentz index, g is the coupling strength of the theory and t^a are the generators of the group $SU(N)$ and A_μ^a is the gauge field. The index a denotes the internal index of the adjoint representation of $SU(N)$ which runs from 1 to $N^2 - 1$.

By definition, the covariant derivative transforms under $SU(N)$ in precisely the same form as $\psi(x)$, namely,

$$\not{D}'\psi(x) \rightarrow U(x)[\not{D}\psi(x)]. \quad (1.1.5)$$

This condition is sufficient to define the transformation properties of the gauge field, A_μ^a ,

$$A_\mu^a t^a \rightarrow U(x)A_\mu^a t^a U^{-1}(x) + \frac{i}{g}(\partial_\mu U(x))U^{-1}(x). \quad (1.1.6)$$

Given the above transformation of the gauge field A_μ^a , it is clear that a mass term of the form,

$$m^2 A_\mu^a A_a^\mu, \quad (1.1.7)$$

cannot be gauge invariant and hence gauge bosons must be massless. It is interesting to note that whilst explicit mass terms for gauge fields in a Lagrangian cannot exist, it is possible to generate a mass term for gauge fields using the Higgs mechanism [18–20]. This will be discussed in detail in section 1.3 in the context of EW symmetry breaking.

$SU(N)$ for $N \geq 2$ is a set of non-Abelian Lie groups. It follows that the generators t^a do not commute and satisfy the following commutation relation,

$$[t^a, t^b] = if^{abc}t^c, \quad (1.1.8)$$

where f^{abc} is known as the structure constant.

A further assumption in Lagrangian mechanics is that *all* terms that satisfy the imposed constraints and gauge symmetries must be included. A trivial term one can

add to the Lagrangian is a mass term for the quark fields,

$$\mathcal{L}_{\text{mass}} = m\bar{\psi}(x)\psi(x), \quad (1.1.9)$$

which in most contexts of this thesis will be presumed to be massless. For sufficiently large scattering energies and for the processes we will be considering this approximation holds very well.

In addition to the above, we can also define the kinetic term for the gauge field,

$$\mathcal{L}_{\text{kinetic}} = -\frac{1}{4}G_{\mu\nu}^a G_a^{\mu\nu}, \quad (1.1.10)$$

where,

$$G_{\mu\nu}^a = \partial_\mu A_\nu^a - \partial_\nu A_\mu^a - gf^{abc}A_\mu^b A_\nu^c. \quad (1.1.11)$$

The physical implications of a non-Abelian gauge theory reside in the final term of Eq. (1.1.11). This term leads to the gluon triple and quartic self-interactions which are not present in Abelian gauge theories such as quantum electrodynamics (QED).

Collecting all of the above terms together leads to the classical Lagrangian,

$$\mathcal{L}_{\text{classical}} = -\frac{1}{4}G_{\mu\nu}^a G_a^{\mu\nu} + \sum_{\text{flavours}} \bar{\psi}(x)(i\not{D} - m)\psi(x), \quad (1.1.12)$$

where we are summing over all active flavours in the theory.

An interesting point to note is that one could add an additional gauge invariant term to the classical Lagrangian,

$$\mathcal{L}_\theta = \frac{\theta g_s^2}{32\pi^2} G_{\mu\nu}^a \tilde{G}_a^{\mu\nu}, \quad (1.1.13)$$

where θ is the coupling strength (with conventional normalisation terms). $\tilde{G}_a^{\mu\nu}$ is the dual field strength tensor,

$$\tilde{G}_a^{\mu\nu} = \frac{1}{2}\epsilon_{\mu\nu\rho\sigma} G_a^{\rho\sigma}. \quad (1.1.14)$$

This term is the equivalent to the ‘ $E \cdot B$ ’ term that arises in the QED Lagrangian. In the case of QCD, this term would give rise to topological CP-violating effects. However, these effects are not observed in Nature implying that θ is consistent with zero. Why this term does not arise is known as the strong CP problem and is an outstanding issue in the theory of QCD.

The classical Lagrangian, as presented above, is not sufficient to perform calcula-

tions perturbatively. Consider the functional integral,

$$\int \mathcal{D}A \exp(iS[A]), \quad (1.1.15)$$

where $S[A]$ is the action for the free gauge field. Using the classical Lagrangian alone, this path integral is divergent given that there are an infinite number of equivalent gauge configurations of the gauge field for each point in x . Ideally one would like to isolate the physical configurations of the gauge field A^μ only once.

Indeed, it is possible to isolate the physical contributions to the functional integral by employing the prescription established by Faddeev and Popov [21]. Consider a functional acting on the gauge field, $\mathcal{G}[\mathbf{A}_\mu(x)]$, which will define a gauge fixing condition. The functional $\delta(\mathcal{G}[\mathbf{A}_\mu(x)])$ isolates gauge configurations corresponding to $\mathcal{G}[\mathbf{A}_\mu(x)] = 0$. Using this functional, we can insert 1 in the definition of the path integral defined in Eq. (1.1.15) in the form [22],

$$1 = \int \mathcal{D}\alpha(x) \delta(\mathcal{G}[\mathbf{A}^\alpha_\mu(x)]) \det \left(\frac{\delta \mathcal{G}[\mathbf{A}^\alpha_\mu(x)]}{\delta \alpha} \right), \quad (1.1.16)$$

where A^α corresponds to the gauge-transformed field,

$$A^\alpha_\mu(x) = A_\mu + \frac{1}{g} D_\mu \alpha(x). \quad (1.1.17)$$

Applying this gauge fixing to a correlation function has two important implications for the Lagrangian [21]. Firstly, this introduces a new gauge fixing term to the Lagrangian of the form,

$$\mathcal{L}_{g.f} = -\frac{1}{\lambda} \text{Tr}(\mathcal{G}[\mathbf{A}_\mu(x)]^2), \quad (1.1.18)$$

where λ parametrises the gauge choice. A commonly used set of gauge fixing terms are given by the covariant gauges,

$$\mathcal{L}_{\text{cov g.f}} = -\frac{1}{2\lambda} (\partial^\mu \mathbf{t} \cdot \mathbf{A}_\mu)^2. \quad (1.1.19)$$

The second implication is that for certain gauge choices in non-Abelian gauge theories, ghost fields can naturally arise [21],

$$\mathcal{L}_{\text{ghost}} = \partial_\mu \boldsymbol{\eta}^{a\dagger} D_{ab}^\mu \boldsymbol{\eta}^b, \quad (1.1.20)$$

where $\boldsymbol{\eta}$ denotes a ghost field. To cancel the unphysical degree of freedom, these fields

must behave as anti-commuting scalars. Such ghost fields cannot be seen as physical particles because they are forbidden by the spin-statistics theorem, however they can contribute within loop diagrams.

It is also possible to construct a gauge choice such that no ghost fields arise,

$$\mathcal{L}_{\text{axial g.f}} = -\frac{1}{\lambda}(n_\mu \mathbf{t} \cdot \mathbf{A}^\mu)^2, \quad (1.1.21)$$

where n^μ is an arbitrary space-time vector. This set of gauge fixing terms is known as axial gauges.

There is a considerable amount of freedom in choosing a gauge when performing calculations. However, in practice it is often easier to calculate in certain gauges than compared to others. For example, the axial gauges have a significantly more complicated gluon propagator than compared to covariant gauge choices. This additional complexity in the gluon propagator is often worse than the necessity to calculate additional ghost diagrams in covariant gauge choices. Nevertheless, axial gauges are useful to demonstrate that ghost fields do not exist in all gauge choices, reaffirming the statement that the ghost fields are indeed unphysical.

1.2 Electroweak interactions

The EW sector of the SM is described by the theory of Glashow-Weinberg-Salam [15–17], whereby a Lagrangian which respects the gauge symmetries $SU(2)_L \times U(1)_Y$ is spontaneously broken by a Higgs mechanism [18–20] into a residual $U(1)_{EM}$ gauge symmetry. In this section we will briefly review the GWS model and precisely the experimental motivations for considering such a theory.

In the 1960's there was a significant amount of experimental data describing weak interactions. Experimental data suggested that the weak interaction, unlike the electromagnetic and strong forces, was parity violating [23]. At low energies, β decay was described very well by the Fermi model [24], where the amplitude for the decay $\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu$ at leading order is described by,

$$i \frac{G_F}{\sqrt{2}} \bar{u}(\nu_\mu) \gamma^\alpha (1 - \gamma_5) u(\mu) \bar{u}(e) \gamma_\alpha (1 - \gamma_5) v(\bar{\nu}_e), \quad (1.2.1)$$

where the Fermi constant G_F is measured experimentally to be [25],

$$G_F = 1.1663787(6) \times 10^{-5} \text{ GeV}^{-2}. \quad (1.2.2)$$

Given the challenge of constructing a complete gauge invariant and renormalisable theory with four-fermion interactions, it was reasonable to assume that the Fermi model was a low energy effective model of a more encompassing theory. Such a theory must involve the exchange of extremely heavy gauge bosons such that the scattering energies of typical experiments at the time were significantly smaller than the mass of the new gauge bosons. This was later confirmed at the Super Proton anti-Proton Synchrontron (SP \bar{P} S) collider with the discovery of the Z and W bosons in 1983 [26–29], which were indeed found to have large masses. The complete description of EW interactions must be able to explain the heavy W and Z gauge bosons and at the same time recover the Fermi model in low energy interactions.

We begin by considering the ‘unbroken’ GWS Lagrangian. Firstly we will define a weak isospin doublet containing a left-handed lepton and its corresponding neutrino,

$$L_L = \begin{pmatrix} \nu_L \\ e_L \end{pmatrix}. \quad (1.2.3)$$

Similarly for the quarks we define the weak isospin doublet,

$$Q_L = \begin{pmatrix} u_L \\ d_L \end{pmatrix}, \quad (1.2.4)$$

where u_L and d_L are the left-handed up-type quark and down-type quark respectively. Similarly to the left-handed fermion content, we also have the analogous right-handed fermion fields (e_R , u_R and d_R) which are singlets under $SU(2)_L$ gauge transformations.

Using the methodology established in section 1.1, this allows us to construct a gauge invariant fermion contribution to the GWS Lagrangian of the form,

$$\begin{aligned} \mathcal{L}_{\text{fermion}} = \sum_j^{\text{generations}} & \left[\bar{L}_L^j \not{D}_L L_L^j + \bar{Q}_L^j \not{D}_L Q_L^j \right. \\ & + \bar{e}_R^j \not{D}_R e_R^j + \bar{u}_R^j \not{D}_R u_R^j \\ & \left. + \bar{d}_R^j \not{D}_R d_R^j \right], \end{aligned} \quad (1.2.5)$$

where \not{D}_L and \not{D}_R denote two distinct covariant derivatives,

$$\not{D}_L = \gamma^\mu \left(\partial_\mu + ig_1 \frac{Y}{2} B_\mu + ig_2 \frac{\tau^i}{2} W_\mu^i \right), \quad (1.2.6)$$

$$\not{D}_R = \gamma^\mu \left(\partial_\mu + ig_1 \frac{Y}{2} B_\mu \right), \quad (1.2.7)$$

where g_1 is the coupling strength of the $U(1)_Y$ gauge field B_μ and Y denotes the hypercharge operator. g_2 is the coupling strength of the $SU(2)_L$ gauge field, τ^i are the set of three generators of $SU(2)_L$ and W_μ^i are the three components of the $SU(2)_L$ gauge field. These are directly analogous to the terms g , t^a and A_μ^a in Eq. (1.1.4). Explicitly the generators τ^i are the 2×2 Pauli matrices given by,

$$\tau^1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \tau^2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \tau^3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \quad (1.2.8)$$

A notable absence from the Lagrangian in Eq. (1.2.5) is the definition of a right-handed neutrino, ν_R . A right-handed neutrino has never been observed in Nature and therefore it is omitted from the GWS Lagrangian.

In addition to hypercharge, we can define another conserved quantum number called weak isospin, T_W , corresponding to $SU(2)_L$ gauge symmetry. Only the left-handed doublets are charged under this gauge symmetry and have $T_W = \frac{1}{2}$ by convention. We can also define the projection of T_W in the W_3^μ direction, T_{W_3} . Here we can use the Pauli matrix τ^3 to determine the value of T_{W_3} for the individual components within the $SU(2)_L$ doublet.

The currents arising from the $SU(2)_L$ gauge symmetry can either be charged currents or neutral currents. Given that after SSB the residual gauge symmetry must be $U(1)_{EM}$ with only neutral currents, we can derive a simple relation between T_{W_3} , Y and the conventional charge Q ,

$$Q = T_{W_3} + \frac{Y}{2}. \quad (1.2.9)$$

Using this information and knowing the charges Q following SSB, we can infer the hypercharge of all the fermions for our unbroken GWS Lagrangian, as shown in Tab. 1.1. As discussed in section 1.1, explicit mass terms for the gauge fields of our Lagrangian would violate gauge invariance. In addition to this, a fermionic mass term split into

| Fermion | T_W | T_{W_3} | Y | Q |
|---------|---------------|----------------|----------------|----------------|
| ν_L | $\frac{1}{2}$ | $+\frac{1}{2}$ | -1 | 0 |
| e_L | $\frac{1}{2}$ | $-\frac{1}{2}$ | -1 | -1 |
| u_L | $\frac{1}{2}$ | $+\frac{1}{2}$ | $+\frac{1}{3}$ | $+\frac{2}{3}$ |
| d_L | $\frac{1}{2}$ | $-\frac{1}{2}$ | $+\frac{1}{3}$ | $-\frac{1}{3}$ |
| e_R | 0 | 0 | -2 | -1 |
| u_R | 0 | 0 | $+\frac{4}{3}$ | $+\frac{2}{3}$ |
| d_R | 0 | 0 | $-\frac{2}{3}$ | $-\frac{1}{3}$ |

Table 1.1: The quantum numbers corresponding to the fermions in the GWS Lagrangian. T_w denotes the weak isospin, T_{w_3} denotes the component of T_w for W_μ^3 . Y is the hypercharge quantum number and Q is the resulting charge following spontaneous symmetry breaking.

its left-handed and right-handed components would be of the form,

$$\mathcal{L}_{\text{mass}} = m[\bar{\psi}(x)_R\psi(x)_L + \bar{\psi}(x)_L\psi(x)_R], \quad (1.2.10)$$

where L and R denote the left-handed and right-handed components of the spinor fields respectively. This is not gauge invariant under $SU(2)_L$ or $U(1)_Y$ gauge transformations. Clearly explicit mass terms cannot be added to the unbroken GWS Lagrangian, hence we need to generate mass terms using a new approach: the Higgs mechanism.

1.3 The Higgs mechanism

The Higgs mechanism uses the concept of SSB to generate mass terms for our gauge fields [18–20]. The basic principle behind SSB is that whilst the Lagrangian respects the full symmetries of a given theory, the *vacuum state* does not necessarily have to. By introducing a new scalar field with a given interaction potential, the full symmetries of the Lagrangian can be spontaneously broken to a residual set. In this section we will review how this is performed in the GWS Lagrangian such that the mass terms of the W and Z bosons can be generated.

Consider an $SU(2)_L$ doublet of complex scalar fields,

$$\Phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix}, \quad (1.3.1)$$

where ϕ^+ and ϕ^0 are the components of the complex scalar field. In terms of real components this has the form,

$$\Phi = \frac{1}{\sqrt{2}} \begin{pmatrix} \phi_1 + i\phi_2 \\ \phi_3 + i\phi_4 \end{pmatrix}. \quad (1.3.2)$$

This doublet has weak isospin $T_W = \frac{1}{2}$ and hypercharge $Y = 1$. It follows from Eq. (1.2.9) that ϕ^+ has an electromagnetic charge $Q = +1$ and ϕ^0 has a charge of $Q = 0$. We can construct a Lagrangian which is invariant under $SU(2)_L \times U(1)_Y$ gauge transformations for this complex scalar field,

$$\mathcal{L}_{\text{scalar}} = (D_\mu^L \Phi)^\dagger (D_\mu^L \Phi) - V(\Phi) + \mathcal{L}_{\text{Yukawa}}, \quad (1.3.3)$$

where D_μ^L is precisely the same covariant derivative as defined in Eq. (1.2.6) and $V(\Phi)$ is the scalar field potential. The Yukawa interaction Lagrangian, $\mathcal{L}_{\text{Yukawa}}$, denotes the gauge invariant interaction vertices between the fermions in the theory and the new scalar particle,

$$\mathcal{L}_{\text{Yukawa}} = - \sum_{i,j=1}^{\text{generations}} [y_{ij}^u Q_L^i \Phi^c u_R^j + y_{ij}^d Q_L^i \Phi d_R^j + y_{ij}^l L_L^i \Phi e_R^j + \text{h.c.}], \quad (1.3.4)$$

where y_{ij}^f denotes the a matrix of Yukawa couplings between the generations and ‘h.c.’ denotes the Hermitian conjugate of the proceeding terms. Φ^c is the charge conjugate of the complex field, defined by,

$$\Phi^c = i\tau^2 \Phi. \quad (1.3.5)$$

The structure of the scalar potential in Eq. (1.3.3) is constrained by gauge invariance and renormalisability to simply,

$$V(\Phi) = \mu^2 (\Phi^\dagger \Phi) - \lambda (\Phi^\dagger \Phi)^2, \quad (1.3.6)$$

where μ^2 and λ are parameters that are yet to be determined. We can however infer that for the potential to be bounded from below, $\lambda < 0$.

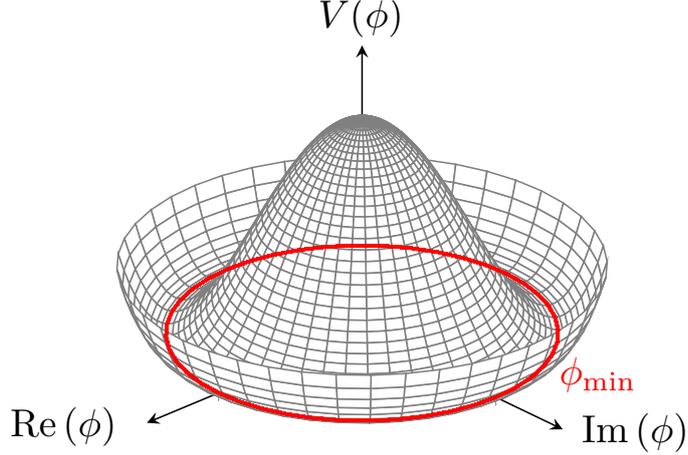


Figure 1.1: The ‘Mexican hat’ Higgs potential for a single complex scalar ϕ , where $V(\phi) = \mu^2 \phi^* \phi - \lambda (\phi^* \phi)^2$ and $\mu^2, \lambda < 0$. The red line denotes the set of field configurations for ϕ which minimises the potential.

If $\mu^2 > 0$ then the minimum of this potential is $V(|\Phi^\dagger \Phi|_{\min}) = 0$ where $|\Phi^\dagger \Phi|_{\min} = 0$. This would imply that μ^2 is a conventional mass term for the scalar field. If however $\mu^2 < 0$, there will be a new set of minima in the potential where $|\Phi^\dagger \Phi|_{\min} \neq 0$ and creates a set of field configurations which reside in the bottom of this potential. This is the famous ‘Mexican hat’ or ‘wine bottle’ Higgs potential, as shown in Fig. 1.1 for a single complex scalar field ϕ . The new minima in $V(\Phi)$ is given by,

$$\frac{dV}{d(\Phi^\dagger \Phi)} = 0 \implies \mu^2 - 2\lambda(\Phi^\dagger \Phi)|_{\min} = 0, \quad (1.3.7)$$

$$\Phi^\dagger \Phi|_{\min} = \frac{\mu^2}{2\lambda}. \quad (1.3.8)$$

We need choose a component in the complex scalar field Φ in which the vacuum expectation value (vev) is non-zero, denoted by v . Doing so will spontaneously break $SU(2)_L \times U(1)_Y$ gauge invariance. In principle this choice is arbitrary, however given that we ultimately wish to describe a theory where a $U(1)_{EM}$ symmetry is preserved, we cannot choose a non-zero vev for either component of the ϕ^+ complex scalar field. This field is charged under $U(1)_{EM}$ gauge transformations, so any non-zero vev will break $U(1)_{EM}$ gauge invariance. Therefore we choose,

$$\langle 0 | \phi_1 | 0 \rangle = \langle 0 | \phi_2 | 0 \rangle = \langle 0 | \phi_4 | 0 \rangle = 0, \quad (1.3.9)$$

$$\langle 0 | \phi_3 | 0 \rangle = v. \quad (1.3.10)$$

From this, consider small perturbations of our field Φ about the vev,

$$\Phi = \frac{1}{\sqrt{2}} \begin{pmatrix} \xi_2 + i\xi_1 \\ v + h + i\xi_3 \end{pmatrix}, \quad (1.3.11)$$

where h denotes the Higgs field and the set of fields $\{\xi_i\}$ are known as Goldstone bosons. It is possible using a suitable $SU(2)_L$ gauge transformation to remove the Goldstone bosons,

$$\Phi = \frac{1}{\sqrt{2}} \exp\left(\frac{i\tau^j \xi^j}{v}\right) \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h \end{pmatrix} + \mathcal{O}(h\xi_i) \rightarrow \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h \end{pmatrix}. \quad (1.3.12)$$

This gauge choice is known as the unitary gauge. Whilst this gauge choice is not particularly useful for performing calculations, it is useful to demonstrate explicitly the mass terms and the particle content of the spontaneously broken Lagrangian. If we further decouple Φ into two components,

$$\Phi = \underbrace{\frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix}}_{\Phi_v} + \underbrace{\frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ h \end{pmatrix}}_{\Phi_h}, \quad (1.3.13)$$

and consider the component Φ_v acting on the complex scalar field Lagrangian in Eq. (1.3.3) and ignore both kinetic and Yukawa terms,

$$\mathcal{L}_{\text{mass}} = \frac{1}{2} \begin{pmatrix} 0 & v \end{pmatrix} \begin{pmatrix} i\frac{g_1}{2}B_\mu + ig_2\frac{\tau^i}{2}W_\mu^i \\ i\frac{g_1}{2}B^\mu + ig_2\frac{\tau^i}{2}W_i^\mu \end{pmatrix} \begin{pmatrix} 0 \\ v \end{pmatrix}. \quad (1.3.14)$$

Expanding out the terms in Eq. (1.3.14) gives,

$$\mathcal{L}_{\text{mass}} = -\frac{v^2}{8} \left[g_2^2 W_\mu^1 W_1^\mu + g_2^2 W_\mu^2 W_2^\mu + (g_1 B^\mu - g_2 W_3^\mu)^2 \right]. \quad (1.3.15)$$

To formulate the mass terms for the spontaneously broken Lagrangian in terms of its mass eigenstates, several substitutions of the gauge fields need to be performed. Firstly, the W_μ^1 and W_μ^2 can be rewritten in terms of W_μ^+ and W_μ^- ,

$$W_\mu^1 = \frac{1}{\sqrt{2}} (W_\mu^+ + W_\mu^-), \quad (1.3.16)$$

$$W_\mu^2 = \frac{i}{\sqrt{2}} (W_\mu^+ - W_\mu^-). \quad (1.3.17)$$

The gauge fields W_μ^3 and B_μ also mix to form the mass eigenstates in a non-trivial manner. Let us define the weak mixing angle, θ_W , which rotates the gauge fields into the mass basis,

$$\begin{pmatrix} A^\mu \\ Z^\mu \end{pmatrix} = \begin{pmatrix} \cos \theta_W & \sin \theta_W \\ -\sin \theta_W & \cos \theta_W \end{pmatrix} \begin{pmatrix} B^\mu \\ W_3^\mu \end{pmatrix}, \quad (1.3.18)$$

where A^μ and Z^μ are defined to be the mass eigenstates. Therefore using the substitutions,

$$A^\mu = \sin \theta_W W_3^\mu + \cos \theta_W B^\mu, \quad (1.3.19)$$

$$Z^\mu = \cos \theta_W W_3^\mu - \sin \theta_W B^\mu, \quad (1.3.20)$$

in conjunction with the definitions of W_μ^1 and W_μ^2 in terms of W_μ^+ and W_μ^- , the Lagrangian for the gauge boson mass terms reads,

$$\begin{aligned} \mathcal{L}_{\text{mass}} = & -\frac{v^2}{4} \left[g_2^2 W_\mu^- W_+^\mu \right. \\ & + Z^\mu Z_\mu (g_1 \sin \theta_W + g_2 \cos \theta_W)^2 + A_\mu A^\mu (g_1 \cos \theta_W - g_2 \sin \theta_W)^2 \\ & \left. - 2Z_\mu A^\mu (g_1 \sin \theta_W + g_2 \cos \theta_W)(g_1 \cos \theta_W - g_2 \sin \theta_W) \right]. \end{aligned} \quad (1.3.21)$$

By definition, there can be no term in the Lagrangian of the form $Z_\mu A^\mu$. This implies that,

$$g_1 \cos \theta_W = g_2 \sin \theta_W, \quad (1.3.22)$$

or alternatively,

$$\sin \theta_W = \frac{g_1}{\sqrt{g_1^2 + g_2^2}}. \quad (1.3.23)$$

Substituting Eq. (1.3.23) into Eq. (1.3.21) reveals that the term of the form $A_\mu A^\mu$ is also zero. This is precisely as we desired; the photon field is massless, meaning the Lagrangian retains a residual $U(1)_{\text{EM}}$ gauge symmetry. The W and Z are the only gauge bosons to have gained a non-zero mass term. Our mass contributions to the Lagrangian simplify to,

$$\mathcal{L}_{\text{mass}} = -\frac{v^2}{4} g_2^2 W_\mu^- W_+^\mu + \frac{v^2}{8} (g_1^2 + g_2^2) Z_\mu Z^\mu, \quad (1.3.24)$$

hence the masses for the W and Z bosons in terms of gauge couplings are,

$$M_W = \frac{vg_2}{2}, \quad M_Z = \frac{vg_2}{2 \cos \theta_W}, \quad (1.3.25)$$

which implies that the weak mixing angle can be defined in terms of the masses of the gauge bosons as,

$$\sin^2 \theta_W = 1 - \frac{M_W^2}{M_Z^2}. \quad (1.3.26)$$

A similar exercise can be performed on the Yukawa interaction terms. Inserting Φ_v into Eq. (1.3.4) reveals,

$$\mathcal{L}_{\text{Yukawa mass}} = \sum_i^{\text{generations}} \left[y_{ii}^u \frac{v}{\sqrt{2}} u_L^i u_R^i + y_{ii}^d \frac{v}{\sqrt{2}} d_L^i d_R^i + y_{ii}^e \frac{v}{\sqrt{2}} e_L^i e_R^i \right], \quad (1.3.27)$$

where for the sake of this discussion we have assumed that the Yukawa coupling matrix is diagonal. This means that the mass terms for a fermion can be read off as,

$$m_f = \frac{y^f v}{\sqrt{2}}. \quad (1.3.28)$$

It is interesting to note that we have not introduced a mass term for the neutrino. If we were to construct a Dirac mass term it would be generated from,

$$\mathcal{L}_{\nu \text{ mass}} = \sum_i^{\text{generations}} y_{ii}^u L_L^i \Phi^c \nu_R^i, \quad (1.3.29)$$

where ν_R would be a right-handed neutrino and a singlet under $SU(2)_L$. This particle has never been observed, however there is now overwhelming evidence that the neutrinos have a small non-zero mass hierarchy. Precisely the origin of this mass remains an unsolved problem in particle physics.

In deriving the mass terms for the spontaneously broken Lagrangian, we neglected the contribution coming from the dynamical higgs field, Φ_h . Repeating this exercise and including Φ_h leads to terms in the Lagrangian from a physical interacting scalar field known as the Higgs boson. The discovery of a scalar particle in run 1 of the LHC [5, 6] which thus far exhibits the properties of the Standard Model Higgs boson is a promising validation of spontaneous symmetry breaking and the Standard Model as a whole.

One of the constraints on the construction of the GWS Lagrangian was that at low energies ($s \ll M_W^2, M_Z^2$) the Fermi model is recovered. Indeed, the tree level amplitude

for the $\nu^- \rightarrow e^- \bar{\nu}_e \nu_\mu$ decay in the unitary gauge for from the GWS theory is given by [30],

$$-i \frac{g_2^2}{2} \bar{u}(\nu_\mu) \gamma^\alpha \left(\frac{1 - \gamma_5}{2} \right) u(\mu) \bar{u}(e) \gamma^\beta \left(\frac{1 - \gamma_5}{2} \right) v(\bar{\nu}_e) \left(-g_{\alpha\beta} + \frac{k_\alpha k_\beta}{M_W^2} \right) \frac{1}{k^2 - M_W^2}. \quad (1.3.30)$$

Which, provided that $k^2 \ll M_W^2$, recovers the prediction of the Fermi model if,

$$\frac{G_F}{\sqrt{2}} = \frac{g_2^2}{8M_W^2} \equiv \frac{e^2}{8 \sin^2 \theta_W M_W^2}. \quad (1.3.31)$$

For the work presented in this thesis, the EW coupling constant α is derived from the relation in Eq. (1.3.31),

$$\alpha \equiv \frac{\sqrt{2} G_F M_W^2}{\pi} \sin^2 \theta_W \equiv \frac{\sqrt{2} G_F M_W^2}{\pi} \left(1 - \frac{M_W^2}{M_Z^2} \right). \quad (1.3.32)$$

This is known as the ‘ G_F -scheme’ or ‘ G_μ -scheme’. This scheme is a choice since higher order electroweak radiative corrections will induce a running of the coupling α . A detailed discussion on scheme choices for α can be found in Ref. [31] in the context of neutral-current EW processes.

1.4 Renormalisation of QCD

It was stated in section 1.1 that all the terms that arise in a physical Lagrangian must be renormalisable. In this section we will review precisely what is meant by a renormalisable theory and why it is necessary to impose this constraint.

To begin, consider a typical generalised tensor integral that would arise when calculating a one-loop diagram in d -dimensions,

$$\mathcal{I}^{\mu_1, \dots, \mu_n}(p_i) = \int \frac{d^d l}{(2\pi)^d} \frac{l^{\mu_1} l^{\mu_2} \dots l^{\mu_n}}{\prod_{i=1}^n [(l^2 - q_i^2) - m_i^2]}, \quad q_i = \sum_{j=0}^i p_j. \quad (1.4.1)$$

Considering just the large l contribution to this integral,

$$\mathcal{I}^{\mu_1, \dots, \mu_n}(p_i) \stackrel{l^2 \gg p_i^2, m_i^2}{\approx} \int \frac{d^d l}{(2\pi)^d} \frac{l^{\mu_1} l^{\mu_2} \dots l^{\mu_n}}{l^{2n}}, \quad (1.4.2)$$

we see that for $n \leq d$ the tensor integral diverges as $l^2 \rightarrow \infty$. This is known as an *Ultraviolet (UV) singularity*. The problem we observe can be understood as a failure in the theory to predict the UV structure of the Lagrangian. The integration of momenta

within a loop can go to arbitrarily high scales where it is presumed that additional terms within a UV complete Lagrangian would contribute. These terms would regulate this UV singularity and generate a finite result.

Given that we cannot reasonably probe the UV structure of Nature, it would appear at first glance that it is not possible to make predictions of quantities on the scale of the interactions without knowing the relevant contributing UV terms. However, a remarkable discovery in QFT [22, 32] states that given a UV complete Lagrangian with terms that appear at a high scale Λ , only the renormalisable contributions to the Lagrangian will not be suppressed by powers of Λ . Therefore, as long as the physical scales in the theory are significantly smaller than Λ , a renormalised Lagrangian is sufficient for collider phenomenology.

The prescription of formulating a renormalised Lagrangian from a so called ‘bare’ Lagrangian requires absorbing divergences within a redefinition of physical parameters in the theory,

$$\psi_0(x) = Z_2^{\frac{1}{2}} \psi(x), \quad (1.4.3)$$

$$A_0^\mu(x) = Z_3^{\frac{1}{2}} A^\mu(x), \quad (1.4.4)$$

$$\eta_0(x) = Z_\eta^{\frac{1}{2}} \eta(x), \quad (1.4.5)$$

$$g_{s,0}^2 = Z_g g_s^2, \quad (1.4.6)$$

$$m_0 = Z_m m, \quad (1.4.7)$$

where terms on the left hand side are the bare parameters, as discussed in section 1.1. Each of the Z_i scaling factors for each measurable quantity are divergent by themselves.

From here, we can define for each Z_i ,

$$Z_i = 1 + \delta_{Z_i}, \quad (1.4.8)$$

which implies that for every contribution to the bare Lagrangian containing a physical parameter, we derive a term proportional to the physical contribution and a term proportional to δ_{Z_i} . The Lagrangian has been decoupled into a renormalised contribution and its counter term,

$$\mathcal{L}_0 = \mathcal{L}_{\text{renorm}} + \mathcal{L}_{\text{c.t.}}, \quad (1.4.9)$$

where the entirety of the UV divergence in \mathcal{L}_0 now resides in $\mathcal{L}_{\text{c.t.}}$.

An important point to note is that the renormalised Lagrangian has lost *some* predictive power. Using the renormalised theory of QCD, it would be impossible to predict the renormalised quantities, such as mass terms and coupling constants, directly from the Lagrangian. However, given a fixed number of measurements, the theory is capable of predicting event rates and differential distributions for a wide range of observables.

To move the UV singularities into the counter term, a method is required to isolate the divergent contributions from the finite contributions, typically using a *regulator*. Many different approaches have been considered in the past [33–35]. For the research relevant for this thesis we will be considering *dimensional regularisation*, where the number of space-time dimensions in the problem is analytically continued from 4-dimensions to d -dimensions, where $d = 4 - 2\epsilon$. The UV divergences then arise as explicit poles in ϵ . Following renormalisation, the physical observables must have a well defined finite limit as $d \rightarrow 4$.

Dimensional regularisation is a convenient regularisation choice given that it respects gauge symmetries and Lorentz invariance. Furthermore, it also simultaneously regulates Infrared (IR) divergences that will be discussed in detail in section 1.6.

Given that the partitioning of the bare Lagrangian into two contributions is somewhat arbitrary, we are free to move finite terms between the renormalised Lagrangian and the counter term. Clearly this has no impact on the Lagrangian as a whole. However, this can have an impact on physical calculations derived from a truncated perturbative series and hence it is important to define the renormalisation scheme when performing perturbative calculations.

The most obvious scheme choice for the finite terms is to define the counter term to have no finite contribution at all and for it to be purely divergent. This is known as the *Minimal Subtraction* (MS) scheme. A popular extension to this scheme is to instead consider additional finite terms that always accompany the pole structure to be part of the counter term. This is called the $\overline{\text{MS}}$ scheme. Explicitly the additional terms are always,

$$\frac{1}{\epsilon} + \ln(4\pi) - \gamma_E + \mathcal{O}(\epsilon), \quad (1.4.10)$$

where γ_E is the Euler-Mascheroni constant (0.5772...). It follows that a rescaling of

the poles by a multiplicative factor,

$$\frac{1}{\epsilon} \rightarrow (4\pi)^\epsilon e^{-\epsilon\gamma_E} \frac{1}{\epsilon} \equiv \bar{C}(\epsilon) \frac{1}{\epsilon}, \quad (1.4.11)$$

is sufficient to move the additional terms into the counter term of the Lagrangian.

A further issue that arises with renormalisation is the choice of *scale* at which we renormalise the theory, μ_R . The renormalisation scale can be considered to be a parametrisation of our ignorance of the scale Λ at which the non-renormalisable UV complete Lagrangian is defined. The choice of μ_R is arbitrary and the physical observables should be independent of this scale choice. If we consider a general observable \mathcal{R} which depends on a single momentum scale s , we can write the independence of the choice of μ_R for this particular observable as,

$$\mu_R^2 \frac{d}{d\mu_R^2} \mathcal{R} \left(\frac{s}{\mu_R^2}, \alpha_s(\mu_R^2) \right) = 0, \quad (1.4.12)$$

where,

$$\alpha_s(\mu_R^2) = \frac{g_s^2(\mu_R^2)}{4\pi}. \quad (1.4.13)$$

Note that given that we have renormalised the coupling constant g_s , the definition of α_s is now a function of the scale choice μ_R . At first glance this might appear to be contradictory, given that g_s cannot depend on any other dimensionful parameters and is dimensionless in 4-dimensions. By performing basic dimensional analysis on the d dimensional Lagrangian we can conclude,

$$[\mathcal{L}] = d, \quad (1.4.14)$$

$$[m\bar{\psi}(x)\psi(x)] = d \rightarrow [\psi(x)] = \frac{d-1}{2}, \quad (1.4.15)$$

$$[(\partial_\nu A_\mu)^2] = d \rightarrow [A^\mu] = \frac{d}{2} - 1, \quad (1.4.16)$$

where the square brackets denotes the dimension of the enclosed term. It follows from considering the quark-gluon vertex, $g_s\bar{\psi}(x)A\psi(x)$, that,

$$[g_s] = 2 - \frac{d}{2} = \epsilon. \quad (1.4.17)$$

The coupling constant has therefore been promoted to a dimensionful parameter in the d -dimensional Lagrangian.

By splitting the derivative in Eq. (1.4.12) and making the substitution $t = \log(s/\mu_R^2)$, we arrive at,

$$\left[-\frac{\partial}{\partial t} + \beta(\alpha_s(\mu_R^2)) \frac{\partial}{\partial \alpha_s(\mu_R^2)} \right] \mathcal{R}(e^t, \alpha_s(\mu_R^2)) = 0, \quad (1.4.18)$$

where $\beta(\alpha_s(\mu_R^2))$ is the QCD β -function. The QCD β -function characterises how $\alpha_s(\mu_R^2)$ changes as a function of μ_R^2 . From Eq. (1.4.18) we can infer that $\beta(\alpha_s(\mu_R^2))$ is defined as,

$$\beta(\alpha_s(\mu_R^2)) = \mu_R^2 \frac{d}{d\mu_R^2} \alpha_s(\mu_R^2). \quad (1.4.19)$$

For sufficiently small values of $\alpha_s(\mu_R^2)$, one can compute the QCD β -function as a perturbative expansion of $\alpha_s(\mu_R^2)$,

$$\frac{\beta(\alpha_s(\mu_R^2))}{2\pi} = \beta_0 \left(\frac{\alpha_s(\mu_R^2)}{2\pi} \right)^2 + \beta_1 \left(\frac{\alpha_s(\mu_R^2)}{2\pi} \right)^3 + \mathcal{O}(\alpha_s^4(\mu_R^2)). \quad (1.4.20)$$

The terms in the β -function can be calculated by considering the variation of Green's functions with respect to μ_R . Given that μ_R only appears in the counter terms of the Lagrangian, this implies they can be calculated by considering the derivative of the counter terms alone to the order in $\alpha_s(\mu_R^2)$ we wish to expand the β -function to. It can be shown that [22],

$$\beta_0 \equiv b_0 N + b_F n_f = \frac{11C_A - 4T_R n_f}{6}, \quad (1.4.21)$$

$$\beta_1 = \frac{17C_A^2 - 10C_A T_R n_f - 6C_F T_R n_f}{6}, \quad (1.4.22)$$

where,

$$C_A = N, \quad (1.4.23)$$

$$C_F = \frac{N^2 - 1}{2N}, \quad (1.4.24)$$

$$T_R = \frac{1}{2}, \quad (1.4.25)$$

for an $SU(N)$ gauge theory. n_f denotes the number of active fermion flavours in our theory. In the context of QCD² $N = 3$ and $n_f = 5$. β_0 and β_1 are renormalisation scheme independent quantities, whereas higher order coefficients of the β -function are

²Typically the top quark is considered to have a large mass and is not a light degree of freedom at LHC scattering energies. Finite top mass effects will be neglected and we will work in the effective theory that the top quark has been integrated out.

dependent on the renormalisation scheme choice. The three-loop [36, 37], four-loop [38, 39] and very recently the five-loop [40] coefficients have all been calculated in the $\overline{\text{MS}}$ scheme.

If we consider only the leading term in Eq. (1.4.20), the differential equation has the following solution,

$$\alpha_s(\mu_R^2) = \frac{\alpha_s(M_Z^2)}{1 + \frac{\beta_0}{2\pi}\alpha_s(M_Z^2)\log\frac{\mu_R^2}{M_Z^2}}, \quad (1.4.26)$$

where M_Z , the mass of the Z boson, is chosen as the reference scale for α_s . This is a typical choice because of the small experimental errors for measurements at this energy scale.

An important point to note is the sign of the β_0 term in the denominator of Eq. (1.4.26). We see that for the values of N and n_f in the theory of QCD, the value for β_0 is positive. The sign of the β_0 term in Eq. (1.4.26) is the origin of a remarkable property of QCD: *asymptotic freedom*. As the energy scale increases, the coupling strength of the theory *decreases*. For Abelian gauge theories such as QED we observe precisely the opposite. It is this property that fuels tightly bound low energy hadronic states and the entirety of nuclear physics.

From Eq. (1.4.26) it can be seen that problems in the perturbative description of QCD occur when,

$$\frac{\beta_0}{2\pi}\alpha_s(M_Z^2)\log\frac{\mu_R^2}{M_Z^2} \sim -1. \quad (1.4.27)$$

This constraint allows us to define the scale Λ_{QCD} , which is the scale at which our perturbative description of strong interactions is no longer valid and partons form bound states. Approximately the value of Λ_{QCD} is 200 MeV [30], meaning for LHC scattering energies of 7-14 TeV we are well in the regime where the perturbative expansion is convergent.

A valid solution of Eq. (1.4.12) is $\mathcal{R}(1, \alpha_s(\mu_R^2))$ [30]. This has the profound impact that the residual scale dependent terms of the observable \mathcal{R} , taken as a perturbative expansion of $\alpha_s(\mu_R^2)$, is entirely described by the $\alpha_s(\mu_R^2)$ scale dependence. Explicitly this implies,

$$\mathcal{R}(1, \alpha_s(\mu_R^2)) = \sum_{i=1}^n r_i \left(\frac{\alpha_s(\mu_R^2)}{2\pi} \right)^i + \mathcal{O}(\alpha_s^{n+1}(\mu_R^2)), \quad (1.4.28)$$

where n is the order in which the series has been truncated in powers of $\alpha_s(\mu_R^2)$.

The set of coefficients r_i do not depend on the renormalisation scale. For sufficiently complicated higher order calculations there can be scale dependent terms in the r_i coefficients that are expected to cancel amongst themselves. This cancellation is often non-trivial and forms a powerful check on a numerical calculation. Section 4.3 will demonstrate that the r_i coefficients are indeed independent of the renormalisation scale for the calculation presented in this thesis.

1.5 Colour decomposition of matrix elements

In principle, the QCD Lagrangian and renormalisation conditions are sufficient to begin constructing QCD amplitudes and matrix elements. Given the complete set of Feynman rules one can construct matrix elements of arbitrary complexity. Indeed, the limiting constraint is the growth in number of diagrams that need to be considered as additional legs or loops are added to the calculation. In this section, we will review a method of organising the matrix element by performing *colour decomposition*, as established in Refs. [41–47], to greatly simplify the calculation of matrix elements³.

Fundamentally a matrix element represents the transition probability density between quantum states. They are summed over all helicity and colour configurations such that,

$$|\mathcal{M}|^2 \equiv \langle \mathcal{M} | \mathcal{M} \rangle = \langle \mathcal{M} | \lambda, c \rangle \langle \lambda, c | \mathcal{M} \rangle, \quad (1.5.1)$$

for all possible helicity states λ and colour configurations c .

It will prove useful to decompose the QCD matrix element according to its colour basis,

$$\langle \lambda, c | \mathcal{M} \rangle = \langle c | \otimes \langle \lambda | \mathcal{M} \rangle, \quad (1.5.2)$$

where $\langle \lambda | \mathcal{M} \rangle$ defines the *partial amplitude*, which is only dependent on kinematical factors. The task of calculating the matrix element has been reduced to calculating the full set of partial amplitudes required for all colour configurations of the specific process being considered. Once the full set of amplitudes has been calculated, the application of colour decomposition will precisely define the combination of partial amplitudes required to recover the full matrix element.

³A pedagogical discussion of tree level and one-loop colour decomposition can be found in Ref. [48].

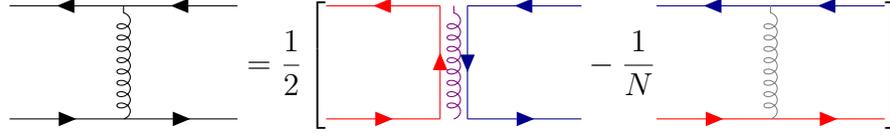


Figure 1.2: A diagrammatic representation of the colour flow in the Fierz identity. In the first term on the right hand side there is a colour exchange between the quark lines. In the second term the gluon does not exchange colour.

The colour structure of a QCD amplitude is dependent on the parton content and the number of loops considered. It is, however, independent of the colourless particles in the process. If a single quark anti-quark pair is in the process then the colour structure can be written in terms of a quark string consisting of multiple gluon emissions between the quark and anti-quark pair,

$$(t^{a_1} t^{a_2} t^{a_3} \dots t^{a_n})_{ij} \equiv t_{ik}^{a_1} t_{kl}^{a_2} t_{lm}^{a_3} \dots t_{mj}^{a_n}. \quad (1.5.3)$$

The Fierz identity can be used to simplify the colour algebra,

$$t_{ij}^a t_{kl}^a = \frac{1}{2} \left(\delta_{il} \delta_{kj} - \frac{1}{N} \delta_{ij} \delta_{kl} \right), \quad (1.5.4)$$

which is derived from the fact that the generators t_{ij}^a form a complete set of traceless hermitian $N \times N$ matrices for an $SU(N)$ gauge theory [22]. The Fierz identity represents the colour flow between two quark anti-quark pairs and a mediating gluon. The term proportional to $\delta_{il} \delta_{kj}$ indicates a colour flow between the two quark pairs, whereas the term proportional to $\delta_{ij} \delta_{kl}$ denotes a colour disconnected contribution. It is this final term that is crucial to the simplification of colour ordered matrix elements. If colour does not flow between the pairs of quarks then the gluon can be considered to not have a charge under $SU(N)$. Since this gluon does not carry a colour charge then it can be considered to be ‘Abelian-like’, implying that non-abelian diagrams will not contribute at this colour level. This will have significant implications for the infrared singularity structure of subleading colour contributions which will be discussed in detail in section 1.6. The commutation relation for the generators of the group in Eq. (1.1.8) can also be applied to decompose factors of f^{abc} ,

$$f^{abc} = -2i(\text{Tr}(t^a t^b t^c) - \text{Tr}(t^a t^c t^b)). \quad (1.5.5)$$

Repeated application of Eqs. (1.5.4) and (1.5.5) are sufficient to colour decompose a QCD process of arbitrary complexity.

1.6 Infrared singularities

In section 1.4 singularities originating from extrapolating the QCD Lagrangian to arbitrarily high scales were discussed. In this section we will review a second classification of divergences arising from parton emissions at arbitrarily *low* scales, known as Infrared (IR) divergences. IR divergences are observed both in loop corrections, where the internal momenta go arbitrarily small, and also in real phase space integrals when additional real emissions are unresolved. In the analyses presented by Bloch and Nordsieck [8]; Kinoshita [9], Lee and Nauenberg [10] (KLN) it was demonstrated that the sum over all degenerate states from real and virtual corrections at a given order in $\alpha_s(\mu_R^2)$ must be finite.

1.6.1 Virtual IR singularities at one-loop

In this section we will review the infrared singular behaviour of one-loop virtual corrections in QCD, as established in Refs. [49–51]. Using the notation in Ref. [49], the universal one-loop pole structure of a QCD amplitude for n parton scattering is of the form,

$$|\mathcal{M}_n^{(1)}(\epsilon; \mu_R^2, \{p_i\})\rangle = \mathbf{I}_n^{(1)}(\epsilon; \mu_R^2, \{p_i\}) |\mathcal{M}_n^{(0)}(\mu_R^2, \{p_i\})\rangle + |\mathcal{M}_{n,\text{finite}}^{(1)}(\mu_R^2, \{p_i\})\rangle, \quad (1.6.1)$$

where $|\mathcal{M}_n^{(1)}(\epsilon; \mu_R^2, \{p_i\})\rangle$ and $|\mathcal{M}_n^{(0)}(\mu_R^2, \{p_i\})\rangle$ are the one-loop and tree level amplitudes respectively. $|\mathcal{M}_{n,\text{finite}}^{(1)}(\mu_R^2, \{p_i\})\rangle$ is finite in the limit of $\epsilon \rightarrow 0$ and process dependent. The entirety of the IR divergences at one-loop resides in the Catani operator $\mathbf{I}_n^{(1)}(\epsilon; \mu_R^2, \{p_i\})$; an operator which is independent of the colourless partons in the process. Explicitly this operator is [49],

$$\mathbf{I}_n^{(1)}(\epsilon; \mu_R^2, \{p_i\}) = \frac{1}{2} \frac{e^{-\epsilon\gamma_E}}{\Gamma(1-\epsilon)} \sum_i^{\text{partons}} \frac{1}{\mathbf{T}_i^2} \nu_i^{\text{sing}}(\epsilon) \sum_{j \neq i} \mathbf{T}_i \cdot \mathbf{T}_j \left(\frac{\mu_R^2 e^{-i\lambda_{ij}\pi}}{2p_i \cdot p_j} \right)^\epsilon, \quad (1.6.2)$$

where $e^{-i\lambda_{ij}\pi}$ denotes a unitarity phase ($\lambda_{ij} = +1$ if i and j are both incoming or outgoing partons and $\lambda_{ij} = 0$ otherwise). The function $\nu_i^{\text{sing}}(\epsilon)$ encapsulates the singularity

structure of the Catani operator,

$$\nu_i^{\text{sing}}(\epsilon) = \mathbf{T}_i^2 \frac{1}{\epsilon^2} + \gamma_i \frac{1}{\epsilon}. \quad (1.6.3)$$

The operators \mathbf{T}_i satisfy the following constraints,

$$\sum_{j \neq i} \mathbf{T}_j = -\mathbf{T}_i, \quad (1.6.4)$$

$$\mathbf{T}_q^2 = \mathbf{T}_{\bar{q}}^2 = C_F, \quad (1.6.5)$$

$$\mathbf{T}_g^2 = C_A, \quad (1.6.6)$$

and the coefficients γ_i depend on the flavour,

$$\gamma_q = \gamma_{\bar{q}} = \frac{3}{2} C_F, \quad (1.6.7)$$

$$\gamma_g = b_0 N + b_F n_f \equiv \beta_0. \quad (1.6.8)$$

The full one-loop matrix is defined as a projection of the one-loop amplitude onto the tree level amplitude,

$$M_n^1 = \langle \mathcal{M}^{(1)} | \mathcal{M}^{(0)} \rangle + \langle \mathcal{M}^{(0)} | \mathcal{M}^{(1)} \rangle \equiv 2 \text{Re} (\langle \mathcal{M}^{(1)} | \mathcal{M}^{(0)} \rangle), \quad (1.6.9)$$

therefore the pole structure of the full matrix one-loop element is given by,

$$\text{Poles}[M_n^1] = 2 \text{Re} (\mathbf{I}_n^{(1)}(\epsilon; \{p_i\})) M_n^0. \quad (1.6.10)$$

When we calculate a one-loop matrix element in d -dimensions, there is a freedom to choose which states are given in 4-dimensions and which are in d -dimensions. The 't Hooft-Veltman (HV) scheme [35] states that all external states are treated in 4-dimensions, whereas the internal states are in d -dimensions. The research in this thesis uses the Conventional Dimensional Regularisation (CDR) scheme exclusively, where *all* of the states are defined in d -dimensions. Any physical observable must be scheme independent, nevertheless it is important that different components of a calculation use a consistent scheme choice. Simple conversions can be made between the scheme choices for one-loop amplitudes which are discussed in Refs. [52–54].

Using the colour decomposition techniques described in section 1.5, the generalised Catani operator in Eq. (1.6.2) can be written as a sum of colour ordered Catani oper-

ators,

$$\mathbf{I}_n^{(1)}(\epsilon; \{p_i\}) = \sum_{ij} \mathbf{I}_{ij}^{(1)}(\epsilon; s_{ij}), \quad (1.6.11)$$

where the summation over ij runs between colour connected partons only. The fact that singularities only occur between neighbouring partons in the colour ordered string is a useful tool when concerned with regulating the IR singularity structures of a one-loop matrix element. Explicitly the colour decomposed operators are [11],

$$\mathbf{I}_{q\bar{q}}^{(1)}(\epsilon; s_{q\bar{q}}) = -\frac{e^{\epsilon\gamma_E}}{2\Gamma(1-\epsilon)} \left[\frac{1}{\epsilon^2} + \frac{3}{2\epsilon} \right] \text{Re}(-s_{q\bar{q}})^{-\epsilon}, \quad (1.6.12)$$

$$\mathbf{I}_{qg}^{(1)}(\epsilon; s_{qg}) = -\frac{e^{\epsilon\gamma_E}}{2\Gamma(1-\epsilon)} \left[\frac{1}{\epsilon^2} + \frac{5}{3\epsilon} \right] \text{Re}(-s_{qg})^{-\epsilon}, \quad (1.6.13)$$

$$\mathbf{I}_{gg}^{(1)}(\epsilon; s_{gg}) = -\frac{e^{\epsilon\gamma_E}}{2\Gamma(1-\epsilon)} \left[\frac{1}{\epsilon^2} + \frac{11}{6\epsilon} \right] \text{Re}(-s_{gg})^{-\epsilon}, \quad (1.6.14)$$

$$\mathbf{I}_{q\bar{q},F}^{(1)}(\epsilon; s_{q\bar{q}}) = 0, \quad (1.6.15)$$

$$\mathbf{I}_{qg,F}^{(1)}(\epsilon; s_{qg}) = -\frac{e^{\epsilon\gamma_E}}{2\Gamma(1-\epsilon)} \frac{1}{6\epsilon} \text{Re}(-s_{qg})^{-\epsilon}, \quad (1.6.16)$$

$$\mathbf{I}_{gg,F}^{(1)}(\epsilon; s_{gg}) = -\frac{e^{\epsilon\gamma_E}}{2\Gamma(1-\epsilon)} \frac{1}{3\epsilon} \text{Re}(-s_{gg})^{-\epsilon}, \quad (1.6.17)$$

where Eqs. (1.6.12), (1.6.13) and (1.6.14) are the leading colour (or leading N) pole contributions. Eqs. (1.6.15), (1.6.16) and (1.6.17) are leading n_f pole contributions.

1.6.2 Virtual IR singularities at two-loops

Using a similar prescription to section 1.6.1, the generalised Catani operators for a two-loop amplitude have also been derived [55],

$$\begin{aligned} |\mathcal{M}_n^{(2)}(\epsilon; \mu_R^2, \{p_i\})\rangle &= \mathbf{I}_n^{(2)}(\epsilon; \mu_R^2, \{p_i\}) |\mathcal{M}_n^{(0)}(\mu_R^2, \{p_i\})\rangle \\ &+ \mathbf{I}_n^{(1)}(\epsilon; \mu_R^2, \{p_i\}) |\mathcal{M}_n^{(1)}(\mu_R^2, \{p_i\})\rangle \\ &+ |\mathcal{M}_{n,\text{finite}}^{(2)}(\mu_R^2, \{p_i\})\rangle, \end{aligned} \quad (1.6.18)$$

where in an analogous format to the one-loop case, the function $|\mathcal{M}_{n,\text{finite}}^{(2)}(\mu_R^2, \{p_i\})\rangle$ is finite as $\epsilon \rightarrow 0$ and is process dependent. The two-loop Catani operator $\mathbf{I}_n^{(2)}(\epsilon; \mu_R^2, \{p_i\})$ has been introduced in a similar fashion to the one-loop Catani operator and is only dependent on the parton content of the process,

$$\mathbf{I}_n^{(2)}(\epsilon; \mu_R^2, \{p_i\}) = -\frac{1}{2} \mathbf{I}_n^{(1)}(\epsilon; \mu_R^2, \{p_i\}) \left(\mathbf{I}_n^{(1)}(\epsilon; \mu_R^2, \{p_i\}) + 4\pi\beta_0 \frac{1}{\epsilon} \right)$$

$$\begin{aligned}
& + \frac{e^{-\epsilon\gamma_E}\Gamma(1-2\epsilon)}{\Gamma(1-\epsilon)} \left(2\pi\beta_0\frac{1}{\epsilon} + K \right) \mathbf{I}_n^{(1)}(2\epsilon; \mu_R^2, \{p_i\}) \\
& + H^{(2)}(\epsilon; \mu_R^2, \{p_i\}), \tag{1.6.19}
\end{aligned}$$

where,

$$K = \left(\frac{67}{18} - \frac{\pi^2}{6} \right) C_A - \frac{10}{9} T_R n_f. \tag{1.6.20}$$

the function $H^{(2)}(\epsilon; \mu_R^2, \{p_i\})$ is a new feature of the generalised pole structure appearing at two-loops. This contains poles at most of order $1/\epsilon$ and is dependent on the parton content and renormalisation scheme choice. It can only be calculated as the remainder between a full two-loop calculation with the relevant parton content and the rest of the Catani pole structure. In practice, these terms can be calculated directly from simple processes such as colourless final state production and applied to more complicated calculations.

The full two-loop matrix element is constructed from a projection of the two-loop amplitude onto the tree level amplitude and the projection of the one-loop amplitude onto itself,

$$M_n^2 = \langle \mathcal{M}_n^{(2)} | \mathcal{M}_n^{(0)} \rangle + \langle \mathcal{M}_n^{(0)} | \mathcal{M}_n^{(2)} \rangle + \langle \mathcal{M}_n^{(1)} | \mathcal{M}_n^{(1)} \rangle, \tag{1.6.21}$$

hence the pole structure for the complete two-loop matrix element is given by,

$$\begin{aligned}
\text{Poles}[M_n^2] &= 2 \mathbf{I}_n^{(1)}(\epsilon; \{p_i\}) \left(M_n^1 - \frac{\beta_0}{\epsilon} M_n^0 \right) \\
& - 2 \left(\mathbf{I}_n^{(1)}(\epsilon; \{p_i\}) \right)^2 M_n^0 \\
& + 2e^{-\epsilon\gamma_E} \frac{\Gamma(1-2\epsilon)}{\Gamma(1-\epsilon)} \left(\frac{\beta_0}{\epsilon} + K \right) \mathbf{I}_n^{(1)}(2\epsilon; \{p_i\}) M_n^0 \\
& + 2H^{(2)}(\epsilon) M_n^0. \tag{1.6.22}
\end{aligned}$$

From Eq. (1.6.22) it is possible to construct a library of colour decomposed Catani operators in a similar fashion to the one-loop Catani operator. The list of operators is lengthy however so it will not be quoted here.

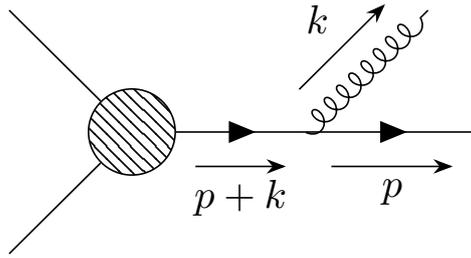


Figure 1.3: A Feynman diagram of a gluon emission from a fermion line. The blob denotes the rest of the hard scattering process.

1.6.3 Real IR singularities at tree level

To demonstrate IR singularities in real phase space integrals, consider the Feynman diagram in Fig. 1.3. The fermion propagator before the gluon emission is given by,

$$\frac{i(\not{p} + \not{k} + m_q)}{(p^\mu + k^\mu)^2 - m_q^2 + i\epsilon}, \quad (1.6.23)$$

where m_q denotes the mass of the quark. If both of the quark and the gluon following emission are on-shell (i.e. $p^\mu p_\mu = k^\mu k_\mu = 0$) and the quark is massless then the propagator simplifies to,

$$\frac{i(\not{p} + \not{k})}{2k^\mu p_\mu + i\epsilon} = \frac{i(\not{p} + \not{k})}{E_k E_p (1 - \cos \theta) + i\epsilon}, \quad (1.6.24)$$

where E_p and E_k are the energies of the emitted quark and gluon respectively and θ is the angle between the emissions.

From Eq. (1.6.24) there are two divergences that can occur,

- soft singularity⁴, $E_k \rightarrow 0$,
- collinear singularity, $\theta \rightarrow 0$.

IR divergences due to unresolved real emissions occur when an intermediate propagator is forced to go on-shell due to a specific kinematical configuration. To extract physical observables from amplitude calculations, it is often necessary to numerically integrate over the entirety of phase space. For this specific propagator it would be possible to regulate the divergence by retaining the mass of the quark in the full amplitude cal-

⁴The additional momentum terms in the numerator of the fermion propagator prevent a divergent contribution for $E_p \rightarrow 0$.

ulation. However, IR divergences would still occur when a massless gluon propagator goes on-shell and hence it is not a general solution to the problem.

In general, the singularity structure of real emissions of a given process factorises into a contribution which is divergent and a matrix element where there is no unresolved emission. For example, the factorisation of a colour ordered tree level matrix element in a soft limit is of the form,

$$M_{n+1}^0(\dots, i, j, k, \dots) \xrightarrow{j \text{ soft}} S_{ijk} M_n^0(\dots, I, K, \dots), \quad (1.6.25)$$

where S_{ijk} is known as an Eikonal factor and is given by,

$$S_{ijk} = \frac{2s_{ik}}{s_{ij}s_{jk}}. \quad (1.6.26)$$

The labels I and K in Eq. (1.6.25) indicate that in the reduced matrix element we have formulated a pair of composite partons from the three partons in the full matrix element. In the strict j soft limit $p_i = p_I$ and $p_k = p_K$. To preserve momentum conservation outside of the strict j soft limit, a smooth mapping function is required to map the momentum configuration p_i, p_j and p_k onto p_I and p_K . The mapping function is a choice and will be discussed in more detail in section 2.1 in the context of antenna subtraction.

In collinear limits, the factorisation is both flavour and helicity configuration dependent. For a pair of final state partons i and j splitting from a final state parton K , the resultant momentum transfer can be parametrised by [22, 49],

$$p_i^\mu = zp^\mu + p_\perp^\mu - \frac{p_\perp^2}{z} \frac{n^\mu}{2p \cdot n}, \quad (1.6.27)$$

$$p_j^\mu = (1-z)p^\mu - p_\perp^\mu - \frac{p_\perp^2}{1-z} \frac{n^\mu}{2p \cdot n}, \quad (1.6.28)$$

where p^μ is a light-like vector defining the collinear direction, z denotes the momentum fraction of parton K transferred to parton i , p_\perp^μ is the momentum component of parton i transverse to the collinear axis and n^μ is an auxiliary light-like vector defined such that $p_\perp \cdot p = p_\perp \cdot n = 0$. p_\perp^μ can be interpreted as a measure of the collinearity of the partons i and j , where in the strict collinear limit $p_\perp^\mu = 0$. The final term proportional to n^μ ensures that the momenta of the particles i and j remain on-shell when p_\perp^μ is non-zero (up to $\mathcal{O}(p_\perp^4)$ corrections).

Unlike in soft limits, the collinear factorisation has a non-trivial spin and flavour dependence,

$$M_{n+1}^0(\dots, i, j, \dots) \xrightarrow{i||j} \frac{\hat{P}_{ij \rightarrow K}^{ss'}}{s_{ij}}(\epsilon; z, p_\perp) M_{n,ss'}^0(\dots, K, \dots), \quad (1.6.29)$$

where s and s' are spin or polarisation indices depending whether parton K is a fermion or vector particle respectively. $\hat{P}_{ij \rightarrow K}^{ss'}$ is a matrix in spinor space and cannot be treated as a multiplicative factor of a reduced matrix element.

In the configuration where K is a fermion, the matrix $\hat{P}_{ij \rightarrow K}^{ss'}$ is proportional to the unity matrix,

$$\hat{P}_{ij \rightarrow K}^{ss'}(\epsilon; z, p_\perp) = \delta_{ss'} P_{ij \rightarrow K}^{(0)}(z), \quad (1.6.30)$$

$P_{ij \rightarrow K}^{(0)}(z)$ denotes the spin averaged splitting functions [56],

$$P_{qq \rightarrow q}^{(0)}(z) = \frac{1 + (1 - z)^2}{z} - \epsilon z, \quad (1.6.31)$$

$$P_{gq \rightarrow q}^{(0)}(z) \equiv P_{qq \rightarrow q}^{(0)}(1 - z) = \frac{1 + z^2}{1 - z} - \epsilon(1 - z), \quad (1.6.32)$$

defined using the CDR scheme.

In a configuration where K is a gluon, the polarisation is labelled by a Lorentz index and is no longer proportional to a unity matrix in spinor space. Explicitly the splitting functions are given by [57],

$$\hat{P}_{q\bar{q} \rightarrow g}^{\mu\nu}(\epsilon; z, p_\perp) = -g^{\mu\nu} + 4z(1 - z) \frac{p_\perp^\mu p_\perp^\nu}{p_\perp^2}, \quad (1.6.33)$$

$$\hat{P}_{gg \rightarrow g}^{\mu\nu}(\epsilon; z, p_\perp) = -2 \left[g^{\mu\nu} \left(\frac{z}{1 - z} + \frac{1 - z}{z} \right) + 2(1 - \epsilon)z(1 - z) \frac{p_\perp^\mu p_\perp^\nu}{p_\perp^2} \right]. \quad (1.6.34)$$

It will prove useful for later discussions to define spin averaged gluon splitting functions,

$$\frac{\hat{P}_{ij \rightarrow g}^{\mu\nu}}{s_{ij}}(\epsilon; z, p_\perp) M_{n,\mu\nu}^0(\dots, K, \dots) = \frac{P_{ij \rightarrow g}^{(0)}(z)}{s_{ij}} M_n^0(\dots, K, \dots) + \text{angular terms}, \quad (1.6.35)$$

where the splitting functions $P_{ij \rightarrow g}^{(0)}(z)$ are taken to be a multiplicative factor of the reduced matrix element instead of being a tensor in spin space, given by [56],

$$P_{q\bar{q} \rightarrow g}^{(0)}(z) = \frac{z^2 + (1 - z)^2 - \epsilon}{1 - \epsilon}, \quad (1.6.36)$$

$$P_{gg \rightarrow g}^{(0)}(z) = \frac{z}{1 - z} + \frac{1 - z}{z} + z(1 - z). \quad (1.6.37)$$

The ‘angular terms’ in Eq. (1.6.35) are a subleading singularity in the collinear limit. It can be shown that following integration over a full phase space these angular terms cancel non-locally between phase space points. Indeed, in cancelling collinear singularities from a phase space point it is sufficient to use the spin averaged splitting functions because of the non-local cancellation of the azimuthal corrections. As a result of this, in future discussions on collinear limits at higher orders in α_s we will always assume spin averaged factorisation. This assumption poses many numerical challenges however which will be discussed in detail in section 4.4.1.

The above factorisation rules for soft and collinear limits can be extended to scenarios involving unresolved emissions from initial state partons. In the case of the soft emissions, this is dealt with automatically through the definition of the invariant, $s_{\hat{i}j}$, where \hat{i} denotes that parton i is in the initial state,

$$s_{\hat{i}j} = (p_j - p_i)^2. \quad (1.6.38)$$

In the case of collinear splitting of an initial state parton \hat{i} into final state partons j and k , the momentum fraction z is taken with respect to the initial state parton,

$$p_j = zp_i + p_\perp^\mu - \frac{p_\perp^2}{z} \frac{n^\mu}{2p \cdot n}, \quad (1.6.39)$$

$$p_k = (1-z)p_i - p_\perp^\mu - \frac{p_\perp^2}{1-z} \frac{n^\mu}{2p \cdot n}, \quad (1.6.40)$$

where similarly to the final state splitting case, p_\perp^μ is the transverse momentum with respect to the collinear axis and n^μ is an auxiliary vector such that $p_i \cdot n = p_\perp \cdot n = 0$. The initial-final spin-averaged splitting functions can be derived from the final state splitting functions [58],

$$P_{q\bar{q} \leftarrow g}^{(0)}(z) = \frac{1}{1-z} \frac{1-\epsilon}{1-z} P_{q\bar{q} \rightarrow g}^{(0)}(z), \quad (1.6.41)$$

$$P_{qg \leftarrow q}^{(0)}(z) = \frac{1}{1-z} P_{qg \rightarrow q}^{(0)}(z), \quad (1.6.42)$$

$$P_{gq \leftarrow q}^{(0)}(z) = \frac{1}{1-z} \frac{1}{1-\epsilon} P_{qg \rightarrow q}^{(0)}(1-z), \quad (1.6.43)$$

$$P_{gg \leftarrow g}^{(0)}(z) = \frac{1}{1-z} P_{gg \rightarrow g}^{(0)}(z). \quad (1.6.44)$$

Eqs. (1.6.41) and (1.6.43) are called flavour changing splitting functions because the initial state hard radiator has changed flavour. In so doing the number of polarisations

in d -dimensions changes using CDR and hence additional factors of $(1 - \epsilon)$ are included.

1.6.4 Real IR singularities at one-loop

The basic principles of soft and collinear factorisation can be extended to one-loop. In section 1.6.1, it was shown that the one-loop matrix element is constructed from an interference between the tree level amplitude and a one-loop amplitude. This implies that when a real emission in a colour ordered one-loop matrix element goes unresolved, the matrix element must factorise such that [59],

$$M_{n+1}^1(\dots, i, j, k, \dots) \xrightarrow{\text{unresolved}^j} X^0(i, j, k) M_n^1(\dots, I, K, \dots) + X^1(\epsilon; i, j, k) M_n^0(\dots, I, K, \dots), \quad (1.6.45)$$

where the divergent singularity structure of the one-loop matrix element M_{n+1}^1 resides entirely in $X^0(i, j, k)$ and $X^1(\epsilon; i, j, k)$. The $X^0(i, j, k)$ function can either be an Eikonal function in the soft limit or a tree level splitting function, as defined in the previous section. This function describes the singularity structure within the tree level amplitude, hence it is the same set of functions as for tree level unresolved emissions.

The function $X^1(\epsilon; i, j, k)$ is a new term required to describe the singularities within a one-loop amplitude, either a one-loop soft term or a one-loop splitting function. These have all been calculated and have been presented in Refs. [59, 60].

1.6.5 Double real IR singularities at tree level

The discussion in section 1.6.3 can be extended to scenarios where two emissions in the final state go unresolved simultaneously. The functions used to describe the singularity structures depend on both the type of unresolved emissions and also the colour configuration of the partons (i.e. where the unresolved emissions appear in the colour ordered string). There are three possible colour configurations for double unresolved emissions,

- Colour connected unresolved emissions - the two unresolved partons are next to each other in the colour ordered string, as depicted in Fig. 1.4;
- Almost colour connected unresolved emissions - the two unresolved partons are

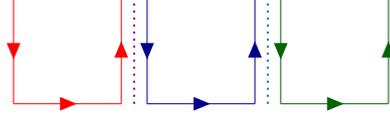


Figure 1.4: Colour connected double unresolved emissions for $q\bar{q} \rightarrow gg$. The vertical dashed lines denote unresolved gluon emissions.

separated by one hard parton in the colour ordered string, as depicted in Fig. 1.5;

- Colour disconnected unresolved emissions - the two unresolved partons are separated by two or more hard partons in the colour ordered string, as depicted in Fig. 1.6.

All possible double unresolved emissions will be reviewed in this section.

Double soft emissions

Double soft emissions occur when two partons go soft simultaneously. This can happen in all colour configurations (colour connected, almost colour connected and colour disconnected) and the functions describing the singularity structure are dependent on the colour connectivity of the partons.

In the case of colour connected gluons j and k going soft, the factorisation of the full matrix element is of the form [61],

$$M_{n+2}^0(\dots, i, j, k, l, \dots) \xrightarrow{j,k \text{ soft}} S_{ijkl} M_n^0(\dots, I, L, \dots), \quad (1.6.46)$$

where as in the single unresolved case partons I and L are composite partons.

Unlike in single soft emissions, a divergence occurs when a pair of same flavour quarks go soft simultaneously. In the colour connected configuration this is analogous to Eq. (1.6.46) with a soft function as found in Ref. [11].

In the case of almost colour connected and colour disconnected configurations, the singular behaviour of the matrix element is described by an iterative implementation of the single soft Eikonal functions,

$$M_{n+2}^0(\dots, i, j, k, \dots, l, m, n \dots) \xrightarrow{j,m \text{ soft}} S_{ijk} S_{lmn} M_n^0(\dots, I, K, \dots, L, N, \dots), \quad (1.6.47)$$

where S_{ijk} and S_{lmn} are precisely the same Eikonal functions as described in section 1.6.3. In the case of almost colour connected configurations then $k = l$ but

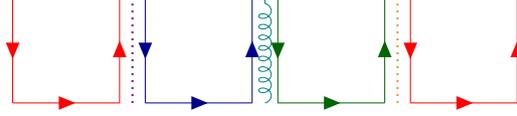


Figure 1.5: Almost colour connected double unresolved emissions for $q\bar{q} \rightarrow ggg$. The vertical dashed lines denote unresolved gluon emissions.

otherwise the singularity structure is the same.

Soft collinear emissions

Soft collinear emissions can occur when a gluon goes soft and a separate pair of partons goes collinear. Similarly to double soft configurations, the divergence depends on the colour separation of the unresolved emissions. If the soft gluon and the collinear pair are colour connected then the factorisation takes the form [62],

$$M_{n+2}^0(\dots, i, j, k, l, \dots) \xrightarrow{j \text{ soft}, k||l} S_{i;jkl} \frac{P_{kl \rightarrow K}^{(0)}(z)}{s_{kl}} M_n^0(\dots, I, K, \dots) + \text{angular terms}, \quad (1.6.48)$$

Similarly to the double soft limits, if the soft gluon and the collinear pair of partons are not colour connected then the singularity is described by an iteration on the single soft Eikonal function and a tree level splitting function,

$$M_{n+2}^0(\dots, i, j, k, \dots, l, m \dots) \xrightarrow{j \text{ soft}, l||m} S_{i;jk} \frac{P_{lm \rightarrow L}^{(0)}(z)}{s_{lm}} M_n^0(\dots, J, L, \dots) + \text{angular terms}. \quad (1.6.49)$$

Double collinear emissions

Double collinear emissions involve two independent pairs of partons going simultaneously collinear. The singularity structure is independent of the separation in colour ordering between the two pairs of partons and is always of the form,

$$M_{n+2}^0(\dots, i, j, \dots, k, l \dots) \xrightarrow{i||j, k||l} \frac{P_{ij \rightarrow I}^{(0)}(z_1)}{s_{ij}} \frac{P_{kl \rightarrow K}^{(0)}(z_2)}{s_{kl}} M_n^0(\dots, I, K, \dots) + \text{angular terms}, \quad (1.6.50)$$

where z_1 and z_2 denote the momentum fractions for the splitting of composite partons I and K respectively.

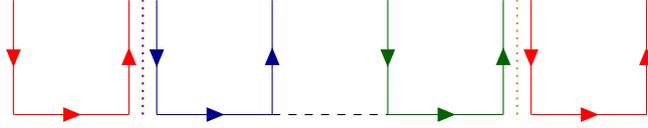


Figure 1.6: Colour disconnected double unresolved emissions for $q\bar{q} \rightarrow ng$. The vertical dashed lines denote unresolved gluon emissions.

Triple collinear emissions

Triple collinear emissions occur in the limit where three partons are all collinear with each other. The singular limit only occurs when all three partons are adjacent in colour ordering [62],

$$M_{n+2}^0(\dots, i, j, k, \dots) \xrightarrow{i||j||k} P_{ijk \rightarrow I}(w, x, y, s_{ij}, s_{jk}, s_{ijk}) M_n^0(\dots, I, \dots) + \text{angular terms}, \quad (1.6.51)$$

where w , x and y denote the momentum fraction of the partons i , j and k with respect to the composite parton I . The splitting function is dependent on the flavour of the partons involved in the triple collinear limit.

1.7 Hadron-hadron collisions and cross sections

Thus far the theory of QCD considered has focused on the concept of free partons at high scattering energies. Given a series of Feynman rules and a set of ‘sensible observables’⁵, the theory of QCD is predictive for lepton-lepton colliders. It is not sufficient for colliders with hadrons in the initial state however, where the strong dynamics of the hadrons cannot be described from first principles in terms of free partons.

The naïve parton model describes hadrons as a collection of free partons carrying a certain momentum fraction z_i of the total momentum of the hadron i [63]. The colliding hadrons are travelling at relativistic speeds in their centre of mass frame, leading to the hadrons being Lorentz contracted in the longitudinal direction. On the scale of the interaction time between the two hadrons, the low energy self-interactions of the hadrons can be neglected and the momentum fractions are definite and fixed.

The probability of finding a given flavour of parton with a specific momentum frac-

⁵The precise definition of a ‘sensible observable’ will be discussed in section 1.8.

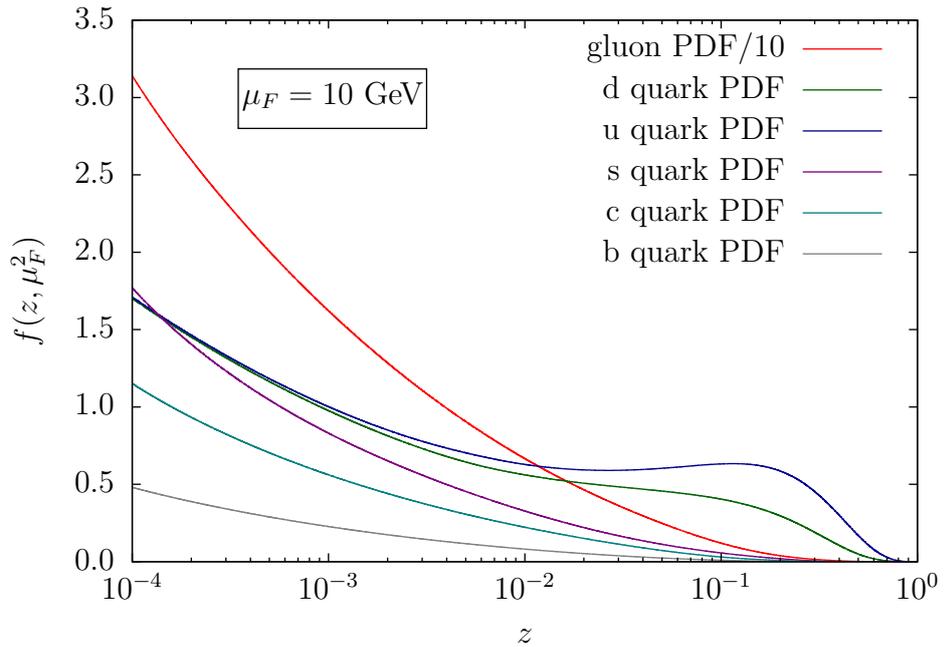


Figure 1.7: An example PDF set from NNPDF2.3 [64] for a proton with $\mu_F = 10$ GeV. The gluon PDF has been reduced by a factor of 10 for clarity.

tion z_i varies as a function of z_i . This leads to the definition of the *parton distribution function* (PDF) which models the short scale physics immediately before the collision as a probabilistic likelihood of finding a parton with a given momentum fraction. PDFs are numerically derived from fitting experimental data and cannot be calculated from first principles. Common features of all PDF sets are that at small momentum fractions the gluon PDFs dominate (diverging as $z \rightarrow 0$) and at large momentum fractions the valence quarks of the hadron dominate.

For a hadron-hadron collider the cross section, which provides a measure for the likelihood of a specific scattering event to occur, can be written in a factorised form [65],

$$d\sigma(P_1, P_2) = \sum_{i,j} \int \frac{dz_1}{z_1} \frac{dz_2}{z_2} f_i(z_1) f_j(z_2) d\hat{\sigma}_{ij}(z_1 P_1, z_2 P_2) + \mathcal{O}(\Lambda_{\text{QCD}}^2/s), \quad (1.7.1)$$

where P_1 and P_2 are the incoming hadron momenta, z_1 and z_2 are the momentum fractions carried by partons i and j contributing to the hard scattering process and $f_i(z_1)$ and $f_j(z_2)$ are their respective PDFs. The term $d\hat{\sigma}_{ij}(z_1 P_1, z_2 P_2)$ denotes the partonic cross section normalised to the hadron-hadron flux and is dependent on the flavour of the colliding partons in the initial state. A summation over all possible partonic combinations, convoluted with their PDFs, returns the hadronic cross section

for a given final state process. Here we neglect terms of order Λ_{QCD}^2/s , where \sqrt{s} is the hadronic center-of-mass energy.

The factorised form of Eq. (1.7.1) provides a foundation on which calculations of hard scattering processes can be calculated in a hadron collider environment. However it is important to note that the validity of factorisation has only been demonstrated for the simplest LHC process, $pp \rightarrow X$, where X is a colourless final state [66]. Given a lack of alternative approaches to computing cross sections however, factorisation is assumed to be a universal property of all hard scattering processes at the LHC.

The partonic cross section can be computed from a perturbative expansion in $\alpha_s(\mu_R^2)$,

$$d\hat{\sigma}_{ij} = \left(\frac{\alpha_s(\mu_R^2)}{2\pi}\right)^n d\hat{\sigma}_{ij}^{\text{LO}} + \left(\frac{\alpha_s(\mu_R^2)}{2\pi}\right)^{n+1} d\hat{\sigma}_{ij}^{\text{NLO}} + \left(\frac{\alpha_s(\mu_R^2)}{2\pi}\right)^{n+2} d\hat{\sigma}_{ij}^{\text{NNLO}} + \mathcal{O}(\alpha_s^{n+3}(\mu_R^2)), \quad (1.7.2)$$

where $d\hat{\sigma}_{ij}^{\text{LO}}$, $d\hat{\sigma}_{ij}^{\text{NLO}}$ and $d\hat{\sigma}_{ij}^{\text{NNLO}}$ are the leading order (LO), next-to-leading order (NLO) and next-to-next-to-leading order (NNLO) contributions to the partonic cross section respectively. n denotes the overall power of $\alpha_s(\mu_R^2)$ within the LO contribution to the calculation.

The partonic cross section for n final state partons for a given perturbative order can be formulated as,

$$d\hat{\sigma}_{ij} = \frac{1}{2s} \sum_{n \geq m} d\Phi_n(p_3, \dots, p_{n+2}; p_1, p_2) \frac{1}{s_n} |\mathcal{M}_{n+2}|^2 J_m^{(n)}(\{p\}_n), \quad (1.7.3)$$

where $d\Phi_n(p_3, \dots, p_{n+2}; p_1, p_2)$ denotes a phase space measure and $|\mathcal{M}_{n+2}|^2$ is the matrix element at a given order in $\alpha_s(\mu_R^2)$ and for a given set of particles. s_n is a symmetry factor for identical partons in the final state. The function $J_m^{(n)}(\{p\}_n)$ is a jet algorithm and provides a mapping from a set of final state partonic momentum $\{p\}_n$ to physical observables. In physical computations it is necessary to impose cuts on experimental observables of final state particles to simulate experimental data. The function $J_m^{(n)}(\{p\}_n)$ can be thought of as a sophisticated ‘Heaviside function’, where if the physical cuts are all satisfied then it has a value of one and is zero otherwise. Jet algorithms will be discussed in further detail in section 1.8.

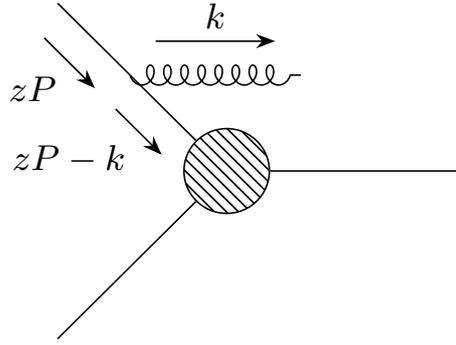


Figure 1.8: Emission of a parton on an initial state leg, changing the instantaneous momentum fraction of the initial state parton.

1.7.1 The improved parton model

The naïve parton model of the previous section factorises the problem of calculating cross sections in terms of a perturbative partonic cross section weighted by a non-perturbative PDF, summed over all flavours of partons. It was assumed that immediately before a hadron-hadron collision, the momentum fraction of the partons within the hadron were frozen and exchanges of soft or unresolved emissions were forbidden.

In describing the partonic cross section, the precision of the calculation was improved by describing it in terms of an expansion of higher order matrix elements where additional emissions can occur and partons can go arbitrarily unresolved in the final state. Similarly the description of our hadronic cross section can be improved by considering additional unresolved *initial* state radiation through a perturbative expansion in α_s .

In the naïve parton model, the factorisation between initial state non-perturbative contributions to the cross section and perturbative hard scattering was clear. Adding initial state unresolved radiation introduces an ambiguity in the distinction between hard scattering processes and the PDFs. Indeed, emissions from the initial state partons can change their instantaneous momentum fraction, in contradiction with the assumption of the previous section that the momentum fraction of the initial state partons is fixed.

In the discussion of renormalisation in section 1.4, ignorance of a UV complete Lagrangian introduced divergences through loop integration. Similarly the ignorance of the distinction between the hard scattering process and the initial state non-perturbative

physics introduces initial state collinear divergences. In Fig. 1.8, if $k||P$ then the resulting phase space integral is divergent.

The prescription established in section 1.4 was to absorb UV divergences into multiplicative factors to define physical quantities from bare quantities. To that end, the initial state collinear singularities can be absorbed into a redefinition of physical PDFs from the ‘bare’ PDFs. Instead of being a multiplicative factor, the divergences are absorbed into a splitting kernel $\Gamma_{ba}(y, \mu_F^2)$ and are convoluted with the bare PDFs $\tilde{f}_b(z)$ [67],

$$f_a(z, \mu_F^2) = \int dx dy \tilde{f}_b(x) \Gamma_{ba}(y, \mu_F^2) \delta(z - xy) \equiv [\tilde{f}_b \otimes \Gamma_{ba}](z, \mu_F^2), \quad (1.7.4)$$

where $\Gamma_{ba}(y, \mu_F^2)$ can be defined as a perturbative expansion in $\alpha_s(\mu_F^2)$,

$$\Gamma_{ba}(y, \mu_F^2) = \delta_{ba} \delta(1 - y) + \left(\frac{\alpha_s(\mu_F^2)}{2\pi} \right) \Gamma_{ba}^1(y, \mu_F^2) + \left(\frac{\alpha_s(\mu_F^2)}{2\pi} \right)^2 \Gamma_{ba}^2(y, \mu_F^2) + \mathcal{O}(\alpha_s^3(\mu_F^2)). \quad (1.7.5)$$

There are several important points to note from Eq. (1.7.5). Firstly a new scale μ_F , known as the *factorisation scale*, has been introduced. In a similar fashion to renormalisation, we have exchanged the ignorance of the distinction between initial state non-perturbative physics and the hard scattering process for a new artificial scale. Any calculation truncated to a given order in α_s^n will retain a residual scale dependence on μ_F formally of the order α_s^{n+1} .

Secondly, the first term in the expansion of the splitting kernel, $\delta_{ba} \delta(1 - y)$, can be observed to be a statement of the naïve parton model. As expected, the improved parton model can be understood as small corrections to the underlying naïve parton model.

To determine the bare PDFs, as found in Eq. (1.7.1), the physical PDFs must be convoluted with the inverse splitting kernel,

$$\tilde{f}_a(z) = \int dx dy f_b(x) \Gamma_{ba}^{-1}(y, \mu_F^2) \delta(z - xy), \quad (1.7.6)$$

where the inverse splitting kernel is given by [67],

$$\Gamma_{ba}^{-1}(y, \mu_F^2) = \delta_{ba} \delta(1 - y) - \left(\frac{\alpha_s(\mu_F^2)}{2\pi} \right) \Gamma_{ba}^1(y, \mu_F^2)$$

$$\begin{aligned}
& - \left(\frac{\alpha_s(\mu_F^2)}{2\pi} \right)^2 \left[\Gamma_{ba}^2(y, \mu_F^2) - \sum_c [\Gamma_{bc}^1 \otimes \Gamma_{ca}^1](z) \right] \\
& + \mathcal{O}(\alpha_s^3(\mu_F^2)).
\end{aligned} \tag{1.7.7}$$

Applying the inverse splitting kernel to the factorised cross section found in Eq. (1.7.1) gives,

$$d\sigma(P_1, P_2) = \sum_{i,j} \int \frac{dz_1}{z_1} \frac{dz_2}{z_2} f_i(z_1, \mu_F^2) f_j(z_2, \mu_F^2) d\hat{\sigma}_{ij}(z_1 P_1, z_2 P_2), \tag{1.7.8}$$

where the partonic cross section has been redefined such that,

$$d\hat{\sigma}_{ij}(z_1 P_1, z_2 P_2) = \int \frac{dx_1}{x_1} \frac{dx_2}{x_2} \Gamma_{ki}^{-1}(x_1, \mu_F^2) \Gamma_{lj}^{-1}(x_2, \mu_F^2) d\tilde{\sigma}_{kl}(x_1 z_1 P_1, x_2 z_2 P_2), \tag{1.7.9}$$

where $\tilde{\sigma}_{kl}(x_1 z_1 P_1, x_2 z_2 P_2)$ is the definition of the parton cross section found in Eq. (1.7.1).

In a practical calculation, the contributions to the inverse splitting considered will be truncated to the same order in α_s as the hard scattering process. To this end, the terms introduced from the inverse splitting kernel to the partonic cross section can be expressed as divergent counter terms at a given order in α_s ,

$$d\hat{\sigma}_{ij}^{\text{LO}}(z_1 P_1, z_2 P_2) = d\tilde{\sigma}_{ij}^{\text{LO}}(z_1 P_1, z_2 P_2), \tag{1.7.10}$$

$$d\hat{\sigma}_{ij}^{\text{NLO}}(z_1 P_1, z_2 P_2) = d\tilde{\sigma}_{ij}^{\text{NLO}}(z_1 P_1, z_2 P_2) + d\hat{\sigma}_{ij,\text{MF}}^{\text{NLO}}(z_1 P_1, z_2 P_2), \tag{1.7.11}$$

$$d\hat{\sigma}_{ij}^{\text{NNLO}}(z_1 P_1, z_2 P_2) = d\tilde{\sigma}_{ij}^{\text{NNLO}}(z_1 P_1, z_2 P_2) + d\hat{\sigma}_{ij,\text{MF}}^{\text{NNLO}}(z_1 P_1, z_2 P_2), \tag{1.7.12}$$

where $d\hat{\sigma}_{ij,\text{MF}}^{\text{NLO}}(z_1 P_1, z_2 P_2)$ and $d\hat{\sigma}_{ij,\text{MF}}^{\text{NNLO}}(z_1 P_1, z_2 P_2)$ denote the divergent counter terms (or mass factorisation terms) at their respective orders in the perturbative series. By substituting in Eq. (1.7.7) the mass factorisation terms can be written explicitly as [67],

$$\begin{aligned}
d\hat{\sigma}_{ij,\text{MF}}^{\text{NLO}}(z_1 P_1, z_2 P_2) &= - \int \frac{dx_1}{x_1} \Gamma_{ki}^1(x_1) d\hat{\sigma}_{kj}^{\text{LO}}(x_1 z_1 P_1, x_2 z_2 P_2), \\
& - \int \frac{dx_2}{x_2} \Gamma_{lj}^1(x_2) d\hat{\sigma}_{il}^{\text{LO}}(x_1 z_1 P_1, x_2 z_2 P_2), \\
d\hat{\sigma}_{ij,\text{MF}}^{\text{NNLO}}(z_1 P_1, z_2 P_2) &= - \int \frac{dx_1}{x_1} \Gamma_{ki}^2(x_1) d\hat{\sigma}_{kj}^{\text{LO}}(x_1 z_1 P_1, x_2 z_2 P_2), \\
& - \int \frac{dx_1}{x_1} \Gamma_{ki}^1(x_1) d\hat{\sigma}_{kj}^{\text{NLO}}(x_1 z_1 P_1, x_2 z_2 P_2), \\
& - \int \frac{dx_2}{x_2} \Gamma_{lj}^2(x_2) d\hat{\sigma}_{il}^{\text{LO}}(x_1 z_1 P_1, x_2 z_2 P_2), \\
& - \int \frac{dx_2}{x_2} \Gamma_{lj}^1(x_2) d\hat{\sigma}_{il}^{\text{NLO}}(x_1 z_1 P_1, x_2 z_2 P_2),
\end{aligned} \tag{1.7.13}$$

$$- \int \frac{dx_1}{x_1} \frac{dx_2}{x_2} \Gamma_{ki}^1(x_1) \Gamma_{lj}^1(x_2) d\hat{\sigma}_{kl}^{\text{LO}}(x_1 z_1 P_1, x_2 z_2 P_2), \quad (1.7.14)$$

noting that contributions proportional to $d\hat{\sigma}_{kj}^{\text{NLO}}$ within $d\hat{\sigma}_{ij,\text{MF}}^{\text{NNLO}}$ consist of both real and virtual corrections, meaning that they occupy multiple phase spaces and contain NLO subtraction terms to regulate the internal IR singularity structure.

Explicitly the coefficients of the splitting kernel are given by [67, 68],

$$\begin{aligned} \Gamma_{ij}^1(x) &= -\frac{1}{\epsilon} P_{ij}^{(0)}(x), \\ \Gamma_{ij}^2(x) &= -\frac{1}{2\epsilon} \left(P_{ij}^{(1)}(x) - \frac{\beta_0}{\epsilon} P_{ij}^{(0)}(x) \right), \end{aligned} \quad (1.7.15)$$

where the one-loop splitting functions $P_{ij}^{(1)}(x)$ are available in Refs. [30, 69].

Given the redefinition of the ‘bare’ PDFs into the physical PDFs, we have all of the constituent ingredients to construct finite cross sections beyond leading order. The only remaining issue is that in formulating the physical PDFs we introduced a dependency on the factorisation scale μ_F . To determine how the PDF varies as a function of μ_F , we exploit the fact that to all orders in $\alpha_s(\mu_F^2)$ the hadronic cross section cannot depend on μ_F . Using a similar approach to the derivation of the renormalisation group equations, one can derive the Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) evolution equation [56, 70, 71],

$$\mu_F^2 \frac{\partial}{\partial \mu_F^2} f_i(x, \mu_F^2) = \frac{\alpha_s}{2\pi} \int_x^1 \frac{d\xi}{\xi} P_{ij} \left(\frac{x}{\xi}, \alpha_s(\mu_F^2) \right) f_j(\xi, \mu_F^2), \quad (1.7.16)$$

where $P_{ij}(x/\xi, \alpha_s(\mu_F^2))$ is an all order splitting function. Expanding in $\alpha_s(\mu_F^2)$ gives,

$$P_{ij} \left(\frac{x}{\xi}, \alpha_s(\mu_F^2) \right) = P_{ij}^{(0)} \left(\frac{x}{\xi} \right) + \frac{\alpha_s(\mu_F^2)}{2\pi} P_{ij}^{(1)} \left(\frac{x}{\xi} \right) + \mathcal{O}(\alpha_s^2(\mu_F^2)). \quad (1.7.17)$$

1.7.2 Regulating IR divergences in cross sections

As discussed in section 1.6, the KLN theorem proves to be a powerful tool in demonstrating higher order perturbative corrections in QCD are indeed finite. As long as we sum over all possible real and virtual corrections at a given order in α_s , this should be sufficient to generate IR finite results for sensible observables.

However, there is a fundamental problem in making this cancellation explicit. The

NLO corrections to a total partonic cross section can be written as,

$$\hat{\sigma}_{ij}^{\text{NLO}} = \int_{d\sigma_{n+1}} d\hat{\sigma}_{ij}^R + \int_{d\sigma_n} \left[d\hat{\sigma}_{ij}^V + d\hat{\sigma}_{ij}^{\text{MF,NLO}} \right], \quad (1.7.18)$$

where R and V denote the real and virtual contributions to the NLO corrections respectively. $d\hat{\sigma}_{ij}^{\text{MF,NLO}}$ are the mass factorisation counter terms as discussed in section 1.7.1. The virtual corrections are integrated over an n -parton phase space, where n is the number of partons in the LO contribution. The real corrections however are integrated over an $(n + 1)$ -parton phase space to accommodate the additional unresolved parton in the final state. It is therefore not immediately obvious how the singularities in the real corrections cancel those from the virtual.

Indeed, the singularity structure of the virtual corrections was shown in section 1.6.1 to be a Laurent expansion in ϵ following dimensional regularisation. The singularity structure of the real corrections is only made explicit following numerical integration over the unresolved region of phase space.

For simple processes, it is possible to write the real phase space integral in d -dimensions and to recover a series of explicit poles in ϵ for the real corrections. As dictated by the KLN theorem, this will precisely cancel the explicit poles from the virtual contribution. The difficulty with this approach occurs for high dimensionality phase space integrals, where the expansion becomes increasingly more difficult. An alternative method would be favourable in approaching this task.

Subtraction

One approach to deal with IR singularities is subtraction. The premise behind subtraction is to reformulate the NLO corrections to a given process as,

$$\begin{aligned} \hat{\sigma}_{ij}^{\text{NLO}} = & \int_{d\sigma_{n+1}} \left[d\hat{\sigma}_{ij}^R - d\hat{\sigma}_{ij}^{S,\text{NLO}} \right] \\ & + \int_{d\sigma_n} \left[d\hat{\sigma}_{ij}^V - d\hat{\sigma}_{ij}^{T,\text{NLO}} \right], \end{aligned} \quad (1.7.19)$$

where each of the square brackets is finite in all unresolved configurations and as $\epsilon \rightarrow 0$. There is an additional constraint that,

$$d\hat{\sigma}_{ij}^{T,\text{NLO}} = - \left[\int_{d\sigma_1} d\hat{\sigma}_{ij}^{S,\text{NLO}} - d\hat{\sigma}_{ij}^{\text{MF,NLO}} \right], \quad (1.7.20)$$

implying that any additional functions or structures we introduce to regulate IR singularities in the real phase space must subsequently be removed from the virtual phase space.

The construction of the subtraction term is such that it must mimic the real matrix element's singular behaviour as it approaches unresolved limits. The finite contribution of the subtraction term is arbitrary, given that whatever finite terms are added to the real phase space integral will be removed in the virtual integration. As a result of this, there is a large degree of freedom in constructing subtraction terms. Ultimately however the final result must be independent of the choice of subtraction scheme, hence it is important when performing large computations that multiple independent computations of the same process are performed using alternative subtraction schemes to validate final results.

The concept of subtraction can be extended to NNLO,

$$\begin{aligned} \hat{\sigma}_{ij}^{\text{NNLO}} &= \int_{d\sigma_{n+2}} \left[d\hat{\sigma}_{ij}^{RR} - d\hat{\sigma}_{ij}^S \right] \\ &+ \int_{d\sigma_{n+1}} \left[d\hat{\sigma}_{ij}^{RV} - d\hat{\sigma}_{ij}^T \right] \\ &+ \int_{d\sigma_n} \left[d\hat{\sigma}_{ij}^{VV} - d\hat{\sigma}_{ij}^U \right], \end{aligned} \quad (1.7.21)$$

where,

$$d\hat{\sigma}_{ij}^U = - \left[\int_{d\sigma_2} d\hat{\sigma}_{ij}^S + \int_{d\sigma_1} \left(d\hat{\sigma}_{ij}^T - d\hat{\sigma}_{ij}^{\text{MF,RV}} \right) - d\hat{\sigma}_{ij}^{\text{MF,VV}} \right]. \quad (1.7.22)$$

The NNLO mass factorisation term $d\hat{\sigma}_{ij,\text{MF}}^{\text{NNLO}}$ described in section 1.7.1 has been divided into two contributions, $d\hat{\sigma}_{ij}^{\text{MF,RV}}$ and $d\hat{\sigma}_{ij}^{\text{MF,VV}}$, depending on the dimensionality of the phase space they occupy.

Slicing

Slicing techniques, first developed in Refs. [72–74], involve a non-local cancellation of divergences by introducing a resolution parameter to the integral. If the kinematics of a phase space are above the resolution parameter then the calculation is performed with one additional parton in the final state. If the kinematics of the phase space imply that the configuration is below the resolution parameter then a second integral

is performed which is IR finite. Explicitly this is,

$$\hat{\sigma}_{ij}^{\text{NLO}} = \int_{d\sigma_{n+1}} d\hat{\sigma}_{ij}^R \Theta(X - X_{\min}) + \int_{d\sigma_{n+1}} d\hat{\sigma}_{ij}^{R,\text{approx}} \Theta(X_{\min} - X) + \int_{d\sigma_n} [d\hat{\sigma}_{ij}^V + d\hat{\sigma}_{ij}^{\text{MF}}], \quad (1.7.23)$$

where X is a parameter dependent on the kinematics and X_{\min} is the resolution parameter. The function $d\hat{\sigma}_{ij}^{R,\text{approx}}$ is an approximation of the real corrections in a divergent configuration that can be integrated analytically and combined with the virtual contribution to render the calculation finite. A powerful check of slicing methods is to make sure that the final results are independent of the resolution parameter X_{\min} .

Indeed, the difficulty facing slicing methods is the fact that ultimately the function $d\hat{\sigma}_{ij}^{R,\text{approx}}$ is only exact in the unresolved limit. This introduces a theoretical systematic error dependent on the X_{\min} parameter introduced. In principle, minimising X_{\min} as far as possible will reduce the systematic error. However, reducing the X_{\min} as far as possible introduces large numerical cancellations between the two integrals. In the specific case of phase space slicing, this introduced large logarithmic cancellations of the form $\log(s_{\min}/\hat{s})$, where s_{\min} is the resolution parameter for phase space slicing. The exercise of performing slicing within a numerical Monte Carlo integrator is reduced to balancing systematic uncertainties against numerical precision. As a result of this at NLO slicing techniques have all but been abandoned, instead preferring local subtraction techniques.

Nevertheless slicing techniques are popular at NNLO with q_T [75] and N-jettiness [76] techniques being used to perform NNLO calculations with up to one coloured particle in the final state. Modern NNLO slicing techniques are based on using a physical observable as opposed to invariants as the resolution parameter. Below the resolution parameter they rely on an approximate resummed calculation to describe the unresolved configurations of the phase space. q_T -subtraction relies on the factorisation theorem of Collins, Soper and Sterman [75] whereas N-jettiness subtraction uses Soft Collinear Effective Theory (SCET) [77]. The appeal of using these techniques in NNLO calculations is that they are significantly simpler to implement than local subtraction. Nevertheless, the balance between numerical precision and systematic dependencies on the resolution parameter remain and are still a challenge for modern NNLO calculations using slicing.

1.8 Jet algorithms and IR safe observables

Jet algorithms, as introduced in section 1.7, provide a mapping from a partonic phase space to one constructed of jets. Experimentally jets are defined to be a highly collimated stream of hadrons whereby the sum of the 4-momenta of the constituent hadrons describes the kinematics of the jet. The translation from final state coloured partons discussed thus far to colourless hadrons seems to require an understanding of non-perturbative hadronisation effects. However, for sufficiently hard jets we can exploit the concept of parton-hadron duality [78, 79], whereby the kinematics of the jet can be described by the underlying partonic hard scattering process.

For a leading order theoretical prediction, there is a one to one correspondence to the number of final state partons and the number of jets in a given process. It follows then that radiative corrections can result in multiple partons clustering to form jets. From this perspective it is clear that higher order corrections give a better description of hard final state jets. The goal in describing a given process should be to include as many orders in α_s as feasibly possible.

To continue further in a comparison between theoretical predictions and experiment, a jet algorithm must be defined. The choice of jet algorithm is important given that it could change the entire description of a final state event. To that end, it will prove useful to define a set of observables that will both be easy to predict using theoretical tools and measure experimentally.

In general, the partonic centre of mass frame is boosted with respect to the hadronic centre of mass frame given that the colliding partons contain a fraction of the momentum of the incoming hadrons. This constrains the set of observables to those which are either invariant under Lorentz boosts or transform simply under boosts. The first such useful observable is the transverse momentum of a parton,

$$p_T = \sqrt{p_x^2 + p_y^2}, \quad (1.8.1)$$

which is simply the quadratic sum of the momentum components that are transverse to the beam axis (traditionally taken to be along the z axis). Clearly this observable is invariant under boosts along the z direction.

Another convenient observable is rapidity, defined by,

$$y = \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right). \quad (1.8.2)$$

Rapidity is a hyperbolic angle which transforms simply under boosts along the z axis. Given a boost along the beam axis, E and p_z transform such that $E' = \gamma(E + \beta p_z)$ and $p'_z = \gamma(p_z + \beta E)$. Therefore the rapidity transforms as,

$$y' = y + \frac{1}{2} \ln \left(\frac{1 + \beta}{1 - \beta} \right), \quad (1.8.3)$$

therefore differences in rapidities between frames are invariant under Lorentz boosts. It is also useful to define the pseudorapidity,

$$\eta = \frac{1}{2} \ln \left(\frac{|p| + p_z}{|p| - p_z} \right), \quad (1.8.4)$$

where in the high energy (massless) limit $\eta \rightarrow y$.

Given a set of simple observables, we can begin to define a jet algorithm. The study of jet algorithms has had a lengthy history with many different algorithms considered. The ‘Snowmass accord’ [80] dictates that a set of desirable properties for a potential jet algorithm are,

- it must be simple to implement both theoretically and experimentally,
- it must be defined to all orders in perturbation theory,
- it must yield finite cross sections at any order in perturbation theory,
- the resulting cross section must be insensitive to hadronisation effects.

It is the third point in this list that will prove to be the most constraining. This statement restricts the set of possible jet algorithms to those which are IR safe. IR safety is the condition that a final state of physical observables must be independent of unresolved emissions in the final state. Two broad types of jet algorithm will be briefly discussed in the following sections: cone algorithms [81, 82] and sequential recombination algorithms [83–87], where examples of both IR safe and IR unsafe jet algorithms will be demonstrated.

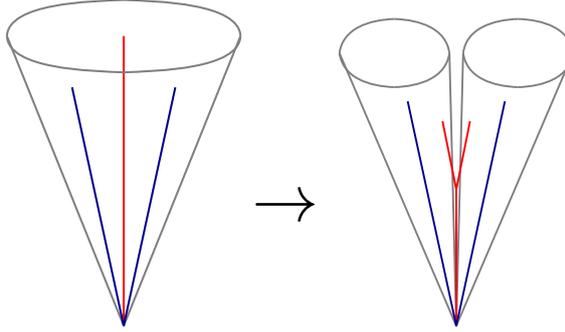


Figure 1.9: A schematic of collinear splitting with seeded cone algorithms that use seeds above a partonic kinematic threshold. A hard parton (red) can undergo collinear splitting into two partons which are both below the kinematic threshold, splitting a larger jet into two smaller jets.

1.8.1 Cone algorithms

Cone algorithms are fundamentally a ‘top-down’ approach to jet clustering. The question asked by the algorithm is whether given a series of final state partons, jets can be fitted to them in a way which is stable to small changes in the configuration?

The first implementations of the cone algorithm were called seeded cone algorithms. The idea being that one would pick a set of momenta as starting direction, typically the hardest parton momenta above a threshold or all of the parton momenta in the final state, as an initial seed. The algorithm would then consider all the partons in a cone of predefined size around the initial seed and sum the 4-momenta of the partons in the cone. This summed momenta will then become the direction for a second iteration of the algorithm. This would be repeated until successive iterations of the algorithm do not change the direction of the cone, i.e. the result becomes stable. The stable condition is such that,

$$D(p_{\text{cone}}, p_a) = 0, \quad (1.8.5)$$

$$p_{\text{cone}} = \sum_i^{\text{all partons}} p_i \Theta(R - D(p_i, p_a)), \quad (1.8.6)$$

where p_a is the reference momentum direction of the cone, R is the jet radius and $D(p_i, p_j)$ is defined to be,

$$D(p_i, p_j) = \sqrt{(y_i - y_j)^2 + (\phi_i - \phi_j)^2}, \quad (1.8.7)$$

where ϕ_i is the azimuthal angle of parton i . Considering the case where our initial

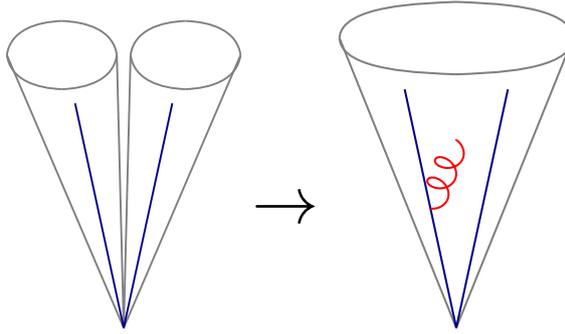


Figure 1.10: A schematic of soft emissions with seeded cone algorithms that use all parton momenta as seeds. A soft emission (red) between two hard partons merges the two jets into one larger jet.

seed is chosen to be all partons above a certain momentum cutoff, the algorithm is not IR safe when undergoing collinear splitting. If a hard parton undergoes collinear splitting, the two daughter partons may fall below the criteria of being classified as ‘hard’. Any scenario where an unresolved emission propagates into a significant change in the topology of the final state is divergent when considering higher order terms in a perturbative expansion. In the case where all parton momenta are considered as seeds, the algorithm is **not IR safe**. A soft emission between two jets causes the two jets to merge into one larger jet, changing the final state topology. Fixes were suggested [88, 89], typically based on an iterative seeding procedure, however even following these fixes it was possible to construct parton configurations that would result in IR unsafe jets. Ultimately seeded cone algorithms were abandoned. The replacement for seeded cone algorithms was to consider all possible starting configurations to find the stable solutions to Eq. (1.8.5), an example being the Seedless Infrared-Safe Cone (SISCone) algorithm [82]. By considering all possible starting positions and finding the solutions it is not possible for soft or collinear emissions to significantly change the stable solutions found.

There are several challenges that remain when considering a seedless cone jet algorithm. Firstly considering all possible seeds to start with potentially leads to a large calculation time for a final state with a large number of particles. The SISCone algorithm uses the geometry of the final state to reduce the number of possible configurations that need to be considered, leading to a good scaling for large final states.

Another issue is to do with overlapping cone structures, typically when two hard jets are found in close proximity. This is particularly challenging to deal with for cone

algorithms because the cone shape has been predefined. The solution to this problem is to consider ‘split-merge’ procedure [81]. The approach either splits the particles in the overlapping region between the two jets according to which they are closest to or merges the two jets. This choice is based on a threshold parameter related to an observable. The split-merge procedure spoils the initial goal of the cone algorithm of fitting perfect cones to particles and leads to distorted event shapes.

1.8.2 Sequential recombination algorithms

Sequential recombination algorithms are complementary to cone algorithms and take a ‘bottom-up’ approach to jet clustering. Instead of attempting to fit cones to the final state, sequential recombination algorithms cluster particles into larger energy deposits. This in essence is an attempt to reverse the fragmentation of QCD emissions to reconstruct the hard scattering process.

To reconstruct jets in hadron-hadron collisions, two distance measures need to be defined,

- d_{ij} - the distance measure between two final state particles i and j ,
- d_{iB} - the distance measure between the final state particle i and the beam.

Given a definition of the distance measures, the clustering of the algorithm is straightforward. For a given final state, compute all of the combinations of d_{ij} and d_{iB} . If the minimum distance measure is d_{ij} , combine particles i and j into a composite particle and restart the algorithm. If d_{iB} is the minimum distance measure, i is considered to be a jet and removed from the set of particles. Repeat the procedure until there are no more particles left and only jets remain.

A popular parametrisation of the distance measures d_{ij} and d_{iB} is given by,

$$d_{ij} = \min \left(p_{T_i}^{2k}, p_{T_j}^{2k} \right) \frac{D(p_i, p_j)^2}{R^2}, \quad (1.8.8)$$

$$d_{iB} = p_{T_i}^{2k}, \quad (1.8.9)$$

where p_{T_i} is the transverse momentum for the particle i , R is the jet radius and the function $D(p_i, p_j)$ is as defined in Eq. (1.8.7). The value of k defines the algorithm being

used; $k = 1$ corresponds to the k_T algorithm [83, 84], $k = 0$ is the Cambridge/Aachen algorithm [85, 86] and $k = -1$ is the anti- k_T algorithm [87].

By construction, sequential recombination algorithms are IR safe. If two partons i and j are collinear then $D(p_i, p_j) \rightarrow 0$, they will therefore be recombined early in the clustering for all values of k . A soft emission cannot significantly impact the combined 4-momentum of a proto-jet after recombination and so soft emissions do not change the final state topology.

The challenges faced by the traditional sequential recombination algorithms are that the jet shape is ill-defined. For the k_T algorithm a distance measure between a hard particle 1 and a soft particle i would be proportional to $p_{T_i}^2$, implying that recombination of partons begins with soft emissions and hard emissions are combined into them. This can result in unusual jet shapes with final states with many soft emissions.

The Cambridge/Aachen algorithm has no dependence on the p_T of the particles in the final state and only discriminates based on angular separation. This naturally forms angularly separated jets but can be heavily distorted by soft emissions.

Finally the anti- k_T algorithm has inverse dependence on the p_T of the particles in its distance measure. This causes jets to cluster around hard and collinear particles in the final state. This has an advantage over other algorithms in that jet shapes are not influenced significantly by soft emissions. This is typical of collider events where there is a hard scattering process combined with an underlying event.

Typically the anti- k_T algorithm produces well defined roughly conical jets when all of the hard partons are well separated angularly. In final states where the jet radii of two hard partons overlap, clearly the two jets cannot both be perfectly conical. The resulting shape is naturally defined by the relative hardness of the two jets, where the harder jet is more conical than the softer jet. In an event with two jets of exactly the same p_T overlap, there will be a straight line separating the two jets. It is for these reasons that the anti- k_T algorithm is favoured in modern experimental analyses and theoretical calculations.

Chapter 2

Antenna Subtraction

As discussed in section 1.6, higher order terms of the perturbative expansion of the partonic cross section have IR singularities that must be regulated to calculate the physical finite result. There is a significant degree of freedom in methods to regulate the divergences from loop integration and phase space integrals and this is an active field of research at NNLO. The work presented in this thesis is based on the antenna subtraction formalism [11–14], a highly flexible subtraction scheme which can be used to regulate divergences in processes with initial and final state QCD radiation up to NNLO in perturbative QCD. In this chapter the basic concepts behind the formalism will be presented, demonstrating examples at NLO and the extensions required at NNLO.

2.1 Unintegrated antenna functions

The antenna subtraction formalism is reliant on the concept of antenna functions. At NLO, the tree level antenna functions are defined such that,

$$X_3^0(i, j, k) = s_{ijk/IK} \frac{M_3^0(i, j, k)}{M_2^0(I, K)}, \quad (2.1.1)$$

where $M_3^0(i, j, k)$ and $M_2^0(I, K)$ are 3 parton and 2 parton colour ordered physical tree level matrix elements respectively, $s_{ijk/IK}$ is a symmetry factor of the final state partons. I and K are composite partons of i, j and k .

The $X_3^0(i, j, k)$ antenna functions are constructed such that they contain *all* of the

tree level singularities that parton j can have between the two hard radiators i and k in this colour ordering. This is, by definition, always true because all QCD matrix elements factorise in the same way¹. The power of the antenna subtraction formalism is that a single function can describe multiple singularities and can smoothly interpolate between them.

Antenna functions can also be defined for NNLO configurations. One-loop singular unresolved antenna, $X_3^1(i, j, k)$, are defined such that [11],

$$X_3^1(i, j, k) = s_{ijk/IK} \frac{M_3^1(i, j, k)}{M_2^0(I, K)} - X_3^0(i, j, k) \frac{M_2^1(I, K)}{M_2^0(I, K)}, \quad (2.1.2)$$

where the divergent contribution from the tree level singularities factorised onto a one-loop matrix element are systematically removed. The library of $X_3^1(i, j, k)$ antennae contain the one-loop unresolved limits appropriate for the flavour and colour configuration under consideration.

Similarly for the colour connected double unresolved emissions, an $X_4^0(i, j, k, l)$ antenna can be defined such that,

$$X_4^0(i, j, k, l) = s_{ijkl,IL} \frac{M_4^0(i, j, k, l)}{M_2^0(I, L)}. \quad (2.1.3)$$

The $X_4^0(i, j, k, l)$ antennae describe the colour connected double resolved tree level configurations. Given a set of libraries of $X_3^0(i, j, k)$, $X_3^1(i, j, k)$ and $X_4^0(i, j, k, l)$ antenna functions, it is possible to construct subtraction terms at NNLO for a process with an arbitrary number of partons in the final state. To describe the full set of singularities possible in a physical process, antennae are required for every possible combination of hard emitters and unresolved partons. To construct such a library, the antenna functions can be categorised on the basis of the possible hard emitters. In QCD there are clearly three hard radiator combinations possible: gg , qg and $q\bar{q}$. The physical matrix elements chosen to form the antenna library must therefore satisfy the following constraints,

- the selected underlying hard radiators must always remain hard. This can be achieved by considering processes involving colourless particles which couple to specific flavours of partons.

¹Neglecting azimuthal terms under collinear splitting which is process dependent.

- The resulting antenna functions must be simple enough that they can be integrated analytically.

It is the second constraint that will prove challenging to satisfy, particularly at NNLO where the analytic integration procedure is significantly more complicated. The three sets of libraries that were chosen are,

- $q\bar{q}$ hard radiators - $\gamma^* \rightarrow q\bar{q} + \text{partons}$ [12]
- gg hard radiators - $H \rightarrow gg + \text{partons}$ [13], where the coupling to the gluons is taken in the infinite top mass effective theory.
- qg hard radiators - $\tilde{\chi} \rightarrow \tilde{g}g + \text{partons}$ [14], where $\tilde{\chi}$ denotes a neutralino and \tilde{g} is a gluino.

The full set of $X_3^0(i, j, k)$, $X_3^1(i, j, k)$ and $X_4^0(i, j, k, l)$ antenna functions are listed in Tab. 2.1. Explicit formulae for the $X_3^0(i, j, k)$ unintegrated antenna functions are given in appendix A. Explicit formulae for all other antennae, including integrated antenna functions, can be found in Ref. [11].

At first glance antenna functions with qg as a pair of hard radiators is unusual. The fundamental issue is that qg does not form a complete colour string, therefore within the Standard Model additional coloured partons must be introduced to complete the colour string. This would complicate the analytical integration of the process and the colour decomposition of the antenna.

A solution to the problem is to appeal to supersymmetry and consider subsequent QCD radiative corrections to a neutralino decay process. The gluino is a fermion, hence it has similar soft and collinear behaviour to a quark. The difference between the quark and the gluino in terms of colour structure resides in the fact that the quark is in the fundamental representation of SU(3) while the gluino is in the adjoint representation. The treatment of their colour algebra is therefore different. Nevertheless, given that the antenna are colour decomposed it is possible to treat the gluino as a quark without issue at NLO. At NNLO complications can occur because the colour structure does not precisely match the desired configuration. Indeed, spurious singularities can appear between colour disconnected particles which must be subtracted from the relevant

$X_4^0(i, j, k, l)$ antenna. The treatment of this issue, amongst other problems related to the neutralino antenna functions, will be discussed in section 2.5.

The $X_3^0(i, j, k)$ and $X_3^1(i, j, k)$ antenna functions are defined as a series of invariants in the 3 parton phase space. Using antenna functions in a general process with an arbitrary number of partons in the final state relies on the condition that the phase space can be factorised such that,

$$d\Phi_{n+1}(p_3, \dots, p_{n+2}; p_1, p_2) = d\Phi_n(p_3, \dots, p_I, p_K, \dots, p_{n+1}; p_1, p_2) d\Phi_{X_{ijk}}(p_i, p_j, p_k) \quad (2.1.4)$$

where $d\Phi_{X_{ijk}}(p_i, p_j, p_k)$ denotes the antenna phase space. The n -parton phase space $d\Phi_n$ is free of singularities when j is unresolved. Similarly for an $X_4^0(i, j, k, l)$ antenna function the phase space factorises such that,

$$d\Phi_{n+2}(p_3, \dots, p_{n+4}; p_1, p_2) = d\Phi_n(p_3, \dots, p_I, p_L, \dots, p_{n+2}; p_1, p_2) d\Phi_{X_{ijkl}}(p_i, p_j, p_k, p_l) \quad (2.1.5)$$

The specific mapping functions used to factorise the phase space depends on whether the hard radiators are in the initial or final state. The mapping functions are not unique, however the possible mapping functions used is heavily constrained by the fact that the mapped partons must remain on-shell and 4-momentum must be conserved.

Final-Final mapping functions

The final-final mapping is appropriate when the two hard radiator partons are both in the final state. A suitable mapping function for $(i, j, k) \rightarrow (I, K)$ is presented in Ref. [90] and reads,

$$p_I^\mu = xp_i^\mu + rp_j^\mu + zp_k^\mu, \quad (2.1.6)$$

$$p_K^\mu = (1-x)p_i^\mu + (1-r)p_j^\mu + (1-z)p_k^\mu, \quad (2.1.7)$$

where,

$$x = \frac{1}{2(s_{ij} + s_{ij})} [(1 + \rho)s_{ijk} - 2rs_{jk}], \quad (2.1.8)$$

| parton content | tree level | one-loop |
|--------------------------------------|--|--|
| <u>$q\bar{q}$ antenna</u> | | |
| $qg\bar{q}$ | $A_3^0(q, g, \bar{q})$ | $A_3^1(q, g, \bar{q}), \tilde{A}_3^1(q, g, \bar{q}), \hat{A}_3^1(q, g, \bar{q})$ |
| $qgg\bar{q}$ | $A_4^0(q, g, g, \bar{q}), \tilde{A}_4^0(q, g, g, \bar{q})$ | |
| $qq'\bar{q}'\bar{q}$ | $B_4^0(q, q', \bar{q}', \bar{q})$ | |
| $qq\bar{q}\bar{q}$ | $C_4^0(q, q, \bar{q}, \bar{q})$ | |
| <u>qg antenna</u> | | |
| qgg | $D_3^0(q, g, g)$ | $D_3^1(q, g, g), \hat{D}_3^1(q, g, g)$ |
| $qggg$ | $D_4^0(q, g, g, g)$ | |
| $qq'\bar{q}'$ | $E_3^0(q, q', \bar{q}')$ | $E_3^1(q, q', \bar{q}'), \tilde{E}_3^1(q, q', \bar{q}'), \hat{E}_3^1(q, q', \bar{q}')$ |
| $qq'\bar{q}'g$ | $E_4^0(q, q', \bar{q}', g), \tilde{E}_4^0(q, q', \bar{q}', g)$ | |
| <u>gg antenna</u> | | |
| ggg | $F_3^0(g, g, g)$ | $F_3^1(g, g, g), \hat{F}_3^1(g, g, g)$ |
| $gggg$ | $F_4^0(g, g, g, g)$ | |
| $gq\bar{q}$ | $G_3^0(g, q, \bar{q})$ | $G_3^1(g, q, \bar{q}), \tilde{G}_3^1(g, q, \bar{q}), \hat{G}_3^1(g, q, \bar{q})$ |
| $gq\bar{q}g$ | $G_4^0(g, q, \bar{q}, g), \tilde{G}_4^0(g, q, \bar{q}, g)$ | |
| $q\bar{q}q'\bar{q}'$ | $H_4^0(q, \bar{q}, q', \bar{q}')$ | |

Table 2.1: A table of all of the tree level and one-loop antenna functions required for NNLO calculations [11]. Antenna denoted by a tilde are subleading colour contributions and antenna denoted with a hat are flavour dependent contributions.

$$z = \frac{1}{2(s_{jk} + s_{ik})} [(1 - \rho)s_{ijk} - 2s_{ij}], \quad (2.1.9)$$

$$\rho^2 = 1 + \frac{4r(1 - r)s_{ij}s_{jk}}{s_{ijk}s_{ik}}, \quad (2.1.10)$$

and r is a free parameter of the mapping function. The value for r chosen is [67],

$$r = \frac{s_{jk}}{s_{ij} + s_{jk}}. \quad (2.1.11)$$

Momentum conservation of the mapping functions is clear from the construction of equations (2.1.6) and (2.1.7). By construction, the mapping function satisfies the

conditions [67],

$$\begin{aligned}
p_I^2 = 0, & & p_K^2 = 0, \\
p_I \rightarrow p_i, & & p_K \rightarrow p_k, & \text{when } j \text{ becomes soft,} \\
p_I \rightarrow p_i + p_j, & & p_K \rightarrow p_k, & \text{when } i \text{ becomes collinear with } j, \\
p_I \rightarrow p_i, & & p_K \rightarrow p_i + p_j, & \text{when } j \text{ becomes collinear with } k.
\end{aligned}$$

The above final-final mapping can be extended to double unresolved emissions between two hard radiators in the final state [91]. However, this mapping is not used in the work presented in this thesis and is only relevant for higher parton multiplicities.

Initial-Final mapping functions

The initial-final mapping function is appropriate when one hard radiator is in the initial state and the other is in the final state. For mapping functions involving an initial state parton, an additional constraint must be satisfied in that the initial state parton must remain on beam axis following the mapping. For the configuration $(\hat{i}, j, k) \rightarrow (\hat{I}, K)$, where \hat{i} and \hat{I} are in the initial state, the mapping function is [58],

$$p_{\hat{I}}^\mu = x p_{\hat{i}}^\mu, \quad (2.1.12)$$

$$p_K^\mu = p_j^\mu + p_k^\mu - (1-x)p_{\hat{i}}^\mu, \quad (2.1.13)$$

where,

$$x = \frac{s_{\hat{i}j} + s_{\hat{i}k} + s_{jk}}{s_{\hat{i}j} + s_{\hat{i}k}}. \quad (2.1.14)$$

Given that x is simply a multiplicative factor $p_{\hat{I}}^\mu$ will remain on beam axis for all possible momentum configurations. The mapping function also satisfies the necessary conditions [58],

$$\begin{aligned}
p_{\hat{I}}^2 = 0, & & p_K^2 = 0, \\
p_{\hat{I}} \rightarrow p_{\hat{i}}, & & p_K \rightarrow p_k, & \text{when } j \text{ becomes soft,} \\
p_{\hat{I}} \rightarrow p_{\hat{i}}, & & p_K \rightarrow p_j + p_k, & \text{when } j \text{ becomes collinear with } k, \\
p_{\hat{I}} \rightarrow p_{\hat{i}} - p_j, & & p_K \rightarrow p_k, & \text{when } j \text{ becomes collinear with } \hat{i}.
\end{aligned}$$

As in the final-final case, the initial-final mapping function can be extended for configurations involving a double unresolved parton emission. For configurations of the form

$(\hat{i}, j, k, l) \rightarrow (\hat{I}, L)$ this is given by an iterated form of the single emission initial-final mapping [58],

$$p_{\hat{I}}^{\mu} = x p_i^{\mu}, \quad (2.1.15)$$

$$p_L^{\mu} = p_j^{\mu} + p_k^{\mu} + p_l^{\mu} - (1-x)p_i^{\mu}, \quad (2.1.16)$$

where,

$$x = \frac{s_{ij} + s_{ik} + s_{il} + s_{jk} + s_{jl} + s_{kl}}{s_{ij} + s_{ik} + s_{il}}. \quad (2.1.17)$$

This function satisfies all of the conditions required in all double unresolved limits [58].

It is interesting to review the behaviour of this mapping function in single unresolved limits. If parton k were to go soft, for example, then the mapping function would be of the form,

$$p_{\hat{I}}^{\mu} \xrightarrow{k \text{ soft}} x p_i^{\mu}, \quad (2.1.18)$$

$$p_L^{\mu} \xrightarrow{k \text{ soft}} p_j^{\mu} + p_l^{\mu} - (1-x)p_i^{\mu}, \quad (2.1.19)$$

where,

$$x = \frac{s_{ij} + s_{il} + s_{jl}}{s_{ij} + s_{il}}, \quad (2.1.20)$$

which is precisely the single unresolved initial-final mapping function given above. In single unresolved limits the phase space must factorise onto the relevant three parton antenna phase space. This is important because the $X_4^0(i, j, k, l)$ antenna functions contain single unresolved limits which factorise onto reduced antenna functions. The double unresolved phase space must factorise in a similar fashion to the double unresolved antenna to regulate the resulting single unresolved divergences.

Initial-Initial mapping functions

The initial-initial mapping functions are appropriate for configurations where both hard radiators are in the initial state. For $(\hat{i}, j, \hat{k}) \rightarrow (\hat{I}, \hat{K})$ the mapping function is [58],

$$p_{\hat{I}}^{\mu} = x p_i^{\mu}, \quad (2.1.21)$$

$$p_{\hat{K}}^{\mu} = y p_k^{\mu}, \quad (2.1.22)$$

$$\tilde{p}_\ell^{\mu} = p_\ell^{\mu} - \frac{2p_\ell \cdot (q + \tilde{q})}{(q + \tilde{q})^2} (q^{\mu} + \tilde{q}^{\mu}) + \frac{2p_\ell \cdot q}{q^2} \tilde{q}^{\mu}, \quad (2.1.23)$$

where ℓ are final state partons that are not involved in the mapping and,

$$q^\mu = p_i^\mu + p_k^\mu - p_j^\mu, \quad (2.1.24)$$

$$\tilde{q}^\mu = xp_i^\mu + yp_k^\mu, \quad (2.1.25)$$

$$x = \sqrt{\frac{s_{\hat{i}\hat{k}} + s_{j\hat{k}}}{s_{\hat{i}\hat{k}} + s_{ij}}} \sqrt{\frac{s_{\hat{i}\hat{k}} + s_{ij} + s_{j\hat{k}}}{s_{\hat{i}\hat{k}}}}, \quad (2.1.26)$$

$$y = \sqrt{\frac{s_{\hat{i}\hat{k}} + s_{ij}}{s_{\hat{i}\hat{k}} + s_{j\hat{k}}}} \sqrt{\frac{s_{\hat{i}\hat{k}} + s_{j\hat{k}} + s_{ij}}{s_{\hat{i}\hat{k}}}}. \quad (2.1.27)$$

To satisfy the condition that both initial state partons must remain on beam axis, the entire reduced phase space configuration must be Lorentz boosted. The mapping function satisfies the necessary conditions [58],

$$\begin{aligned} p_{\hat{I}}^2 &= 0, & p_{\hat{K}}^2 &= 0, \\ p_{\hat{I}} &\rightarrow p_i, & p_{\hat{K}} &\rightarrow p_{\hat{k}}, & \text{when } j \text{ becomes soft,} \\ p_{\hat{I}} &\rightarrow (1 - z_i)p_i, & p_{\hat{K}} &\rightarrow p_{\hat{k}}, & \text{when } j \text{ becomes collinear with } \hat{i}, \\ p_{\hat{I}} &\rightarrow p_i, & p_{\hat{K}} &\rightarrow (1 - z_k)p_{\hat{k}}, & \text{when } j \text{ becomes collinear with } \hat{k}, \end{aligned}$$

and in all unresolved limits $\tilde{p}_\ell \rightarrow p_\ell$.

The extension of the single unresolved initial-initial mapping to a double unresolved mapping is straightforward. For the mapping $(\hat{i}, j, k, \hat{l}) \rightarrow (\hat{I}, \hat{L})$ the mapping is simply [58],

$$p_{\hat{I}}^\mu = xp_i^\mu, \quad (2.1.28)$$

$$p_{\hat{L}}^\mu = yp_l^\mu, \quad (2.1.29)$$

$$\tilde{p}_m^\mu = p_m^\mu - \frac{2p_m \cdot (q + \tilde{q})}{(q + \tilde{q})^2} (q^\mu + \tilde{q}^\mu) + \frac{2p_m \cdot q}{q^2} \tilde{q}^\mu, \quad (2.1.30)$$

where m are final state partons that are not involved in the mapping and,

$$q^\mu = p_i^\mu + p_l^\mu - p_j^\mu - p_k^\mu, \quad (2.1.31)$$

$$\tilde{q}^\mu = xp_i^\mu + yp_l^\mu, \quad (2.1.32)$$

and,

$$x = \sqrt{\frac{s_{\hat{i}\hat{l}} + s_{ij} + s_{lk}}{s_{\hat{i}\hat{l}}}} \sqrt{\frac{s_{\hat{i}\hat{l}} + s_{ij} + s_{ik} + s_{lj} + s_{lk} + s_{jk}}{s_{\hat{i}\hat{l}} + s_{ij} + s_{ik}}}, \quad (2.1.33)$$

| Desired limits | Undesired limits |
|----------------------------|----------------------------|
| 3_g soft | 4_g soft |
| 1_q collinear with 3_g | 1_q collinear with 4_g |
| 3_g collinear with 4_g | |

Table 2.2: A table of the desired and undesired limits of $D_3^0(1_q, 3_g, 4_g)$. It is important that in our antenna 1_q and 4_g act as the hard radiators and do not produce singularities as they go unresolved.

$$y = \sqrt{\frac{s_{i\hat{l}} + s_{ij} + s_{ik}}{s_{i\hat{l}}}} \sqrt{\frac{s_{i\hat{l}} + s_{ij} + s_{ik} + s_{\hat{l}j} + s_{\hat{l}k} + s_{jk}}{s_{i\hat{l}} + s_{ij} + s_{ik}}}. \quad (2.1.34)$$

2.1.1 Sub-antennae

Given the set of antennae defined in Tab. 2.1, we have all of the required components to begin constructing subtraction terms. However, in some cases they are not in precisely the form we desire. Consider the antenna,

$$D_3^0(1_q, 3_g, 4_g) = \frac{1}{s_{134}^2} \left(\frac{2s_{134}^2 s_{14}}{s_{13} s_{34}} + \frac{2s_{134}^2 s_{13}}{s_{14} s_{34}} + \frac{s_{14} s_{34} + s_{34}^2}{s_{13}} \right. \\ \left. + \frac{s_{13} s_{34} + s_{34}^2}{s_{14}} + \frac{2s_{13} s_{14}}{s_{34}} + 5s_{134} + s_{34} \right) + \mathcal{O}(\epsilon). \quad (2.1.35)$$

This antenna contains the divergences described in Tab. 2.2, where they have been divided into limits that should be preserved in the colour ordered antenna and those which should be removed. The cyclic definition of the colour ordering of the neutralino matrix element leads to a colour connection between 1_q and 4_g , meaning that there are singular limits in this antenna where one of the hard radiators becomes unresolved.

There are two approaches to solving this problem. Firstly one could use additional antennae to subtract the spurious singularities. For example introducing structures of the form,

$$D_3^0(1_q, 3_g, 4_g) + D_3^0(5_{\bar{q}}, 3_g, 4_g) - A_3^0(1_q, 4_g, 5_{\bar{q}}), \quad (2.1.36)$$

where the $A_3^0(1_q, 4_g, 5_{\bar{q}})$ antenna function subtracts the unwanted collinear limits $1_q || 4_g$ and $5_{\bar{q}} || 4_g$. The soft limits of such a structure would not be correct however and requires the introduction of additional explicit Eikonal functions to cancel the problematic soft limits.

Clearly this is unwieldy and for lengthy subtraction terms this would be difficult to

manage. A better solution is to decompose the full antenna into two sub-antennae,

$$D_3^0(1_q, 3_g, 4_g) = d_3^0(1_q, 3_g, 4_g) + d_3^0(1_q, 4_g, 3_g), \quad (2.1.37)$$

where there are no singular limits in $d_3^0(1_q, i_g, j_g)$ associated with the j soft configuration. This division can be performed by using partial fractioning on divergent terms proportional to s_{14} in the denominator. Explicitly the final result is [11],

$$d_3^0(1_q, 3_g, 4_g) = \frac{1}{s_{134}^2} \left(\frac{2s_{134}^2 s_{14}}{s_{13} s_{34}} + \frac{s_{14} s_{34} + s_{34}^2}{s_{13}} + \frac{s_{13} s_{14}}{s_{34}} + \frac{5}{2} s_{134} + \frac{1}{2} s_{34} \right) + \mathcal{O}(\epsilon). \quad (2.1.38)$$

In the above sub-antenna it is clear that the 4_g soft limit and the $1_q||4_g$ collinear limit have both been removed. There is some ambiguity on what happens to the $3_g||4_g$ collinear limit however. Performing the substitutions,

$$p_3 \rightarrow z p_{\widehat{(34)}}, \quad (2.1.39)$$

$$p_4 \rightarrow (1-z) p_{\widehat{(34)}}, \quad (2.1.40)$$

where $p_{\widehat{(34)}}$ is the momentum of the composite parton from which partons 3 and 4 split, the $d_3^0(1_q, 3_g, 4_g)$ sub-antenna reduces to,

$$d_3^0(1_q, 3_g, 4_g) \xrightarrow{3_g||4_g} \frac{2}{s_{34}} \left(\frac{(1-z)}{z} + \frac{z(1-z)}{2} \right), \quad (2.1.41)$$

where we have some of the terms in the spin-averaged gluon splitting function in Eq. (1.6.37). The full spin-averaged splitting function is recovered when this antenna is used in conjunction with another equivalent sub-antenna with 3_g and 4_g exchanged².

This sub-antenna is useful for configurations where both gluons, 3_g and 4_g , are in the final state. This cannot be used for processes where one of the gluons is in the initial state and the quark is in the final state however. For example,

$$D_3^0(3_q, 4_g, \hat{1}_g) = d_3^0(3_q, \hat{1}_g, 4_g) + d_3^0(3_q, 4_g, \hat{1}_g), \quad (2.1.42)$$

where $\hat{1}_g$ is the gluon in the initial state. The sub-antenna $d_3^0(3_q, \hat{1}_g, 4_g)$ contains a flavour changing limit $3_q||\hat{1}_g$, where the initial state gluon will become a quark in the

²This can also be an $f_3^0(3_g, 4_g, 5_g)$ sub-antenna which has been constructed to isolate the 4_g soft limit and separates the gluon splitting function in precisely the same way as the $d_3^0(1_q, 4_g, 3_g)$ sub-antenna.

reduced matrix element. At the same time however this sub-antenna contains a part of the spin-averaged gluon splitting function, as outlined above, which is not a flavour changing limit. The symmetric decomposition of the limits in this antenna creates a mismatch in flavour changing configurations.

The solution to this problem is to reformulate the $D_3^0(3_q, 4_g, \hat{1}_g)$ into two new sub-antenna,

$$D_3^0(3_q, 4_g, \hat{1}_g) = d_{3,g \rightarrow q}^0(3_q, \hat{1}_g, 4_g) + d_{3,g \rightarrow g}^0(3_q, 4_g, \hat{1}_g), \quad (2.1.43)$$

where $d_{3,g \rightarrow q}^0(3_q, \hat{1}_g, 4_g)$ contains only the $3_q || \hat{1}_g$ flavour changing collinear limit and $d_{3,g \rightarrow g}^0(3_q, 4_g, \hat{1}_g)$ contains *all* of the remaining flavour preserving limits, including the full spin-averaged gluon collinear limit. The explicit formulae for these antenna functions are given in appendix A.

Ideally the above concepts could be applied to NNLO antenna functions to construct sub-antennae with precisely the correct limits for the appropriate initial state configuration. In some configurations this is possible, for example the final-final F_4^0 antenna can be decomposed successfully into sub-antenna [92]. Unfortunately there are cases where this is not always possible. A particularly problematic antenna is $D_4^0(3_q, \hat{1}_g, 4_g, 5_g)$ which contains many flavour changing and flavour preserving limits with no elegant solution to construct the desired sub-antenna that can be integrated analytically. To use this antenna in a physical process additional antenna structures are required to cancel the problematic flavour changing limits. This will be discussed in detail in section 2.5.

2.2 Integrated antenna functions

As discussed in section 2.1, an important constraint on the possible processes that could be used to construct the antenna libraries is that they are analytically integrable. Great care must be taken in deciding which physical processes will be used to calculate the antenna. In this section the general structure of the integration will be discussed. The precise technical details of performing the analytical integrals can be found in Refs. [93–97].

Given that the phase space factorisation is dependent on the initial state configu-

ration of the hard radiators in an antenna, there are different analytical integrals to be performed in each case. For the final-final three parton antenna functions, the integral is given by [11],

$$\mathcal{X}_3^\ell(s_{ijk}) = \frac{1}{C(\epsilon)} \int d\Phi_{X_{ijk}} X_3^\ell(i, j, k), \quad (2.2.1)$$

where,

$$C(\epsilon) = \frac{\bar{C}(\epsilon)}{8\pi^2}, \quad (2.2.2)$$

and $\bar{C}(\epsilon)$ is defined in Eq. (1.4.11) in the $\overline{\text{MS}}$ scheme. $\ell = 0, 1$, corresponding to a tree level or a one-loop antenna function respectively. The integral of the four parton final-final antenna functions is also defined with a similar relation [11],

$$\mathcal{X}_4^0(s_{ijkl}) = \frac{1}{[C(\epsilon)]^2} \int d\Phi_{X_{ijkl}} X_4^0(i, j, k, l). \quad (2.2.3)$$

In the initial-final phase space mapping, the initial state hard radiator is rescaled $p_i^\mu \rightarrow \hat{x}_i p_i^\mu$, where \hat{x} is dependent on whether it is a single or double parton phase space mapping. Similarly the integral for the initial-final antenna must also be rescaled [95],

$$\mathcal{X}_3^\ell(s_{ijk}; x_1) = \frac{1}{C(\epsilon)} \int d\Phi_2 \delta(x_1 - \hat{x}_1) \frac{Q^2}{2\pi} X_3^\ell(\hat{1}, j, k), \quad (2.2.4)$$

where $Q^2 = -(p_j + p_k - p_1)^2$ and x_1 is the momentum fraction carried by parton $\hat{1}$. The integral is performed over a two parton phase space. The extension to the four parton antenna integral is [95],

$$\mathcal{X}_4^0(s_{ijkl}; x_1) = \frac{1}{[C(\epsilon)]^2} \int d\Phi_3 \delta(x_1 - \hat{x}_1) \frac{Q^2}{2\pi} X_4^0(\hat{1}, j, k, l). \quad (2.2.5)$$

In the initial-initial mapping, both of the initial state partons are rescaled. The resulting integral is given by,

$$\mathcal{X}_3^\ell(s_{\hat{1}j\hat{2}}; x_1, x_2) = \frac{1}{C(\epsilon)} \int dp_j x_1 x_2 \delta(x_1 - \hat{x}_1) \delta(x_2 - \hat{x}_2) X_3^\ell(\hat{1}, j, \hat{2}), \quad (2.2.6)$$

where the extension to this for the four parton antenna function integral is [94],

$$\mathcal{X}_4^0(s_{\hat{1}jk\hat{2}}; x_1, x_2) = \frac{1}{[C(\epsilon)]^2} \int dp_j dp_k x_1 x_2 \delta(x_1 - \hat{x}_1) \delta(x_2 - \hat{x}_2) X_4^0(\hat{1}, j, k, \hat{2}). \quad (2.2.7)$$

For initial-initial and initial-final antenna functions, the final integrated result is dependent on precisely which partons in the antenna are in the initial state. To construct

the complete integrated antenna library, every possible initial state configuration must be considered.

2.3 Antenna subtraction at NLO

The construction of NLO subtraction terms has been fully automated using Catani-Seymour (CS) dipoles [98] and Frixione-Kunszt-Signer (FKS) subtraction [99], hence the structure of NLO subtraction terms is no longer an active field of research. Nevertheless, it will prove to be informative to first review implementing antenna subtraction at NLO and then extending this to NNLO in later sections.

2.3.1 Real subtraction

As discussed in section 1.7.2, the real and virtual corrections to a process must be dealt with independently using subtraction since they are integrated over phase spaces of different dimensionality. The real corrections are of the form,

$$\int_{d\sigma_{n+1}} \left[d\hat{\sigma}_{ij}^R - d\hat{\sigma}_{ij}^{S,\text{NLO}} \right], \quad (2.3.1)$$

where [68],

$$\begin{aligned} d\hat{\sigma}_{ij}^R &= \mathcal{N}_{\text{NLO}}^R \sum_{\mathcal{C}} \mathcal{C}(N, n_f) \sum_{\text{perms}} d\Phi_{n+1}(p_3, \dots, p_{n+3}; p_1, p_2) \frac{1}{S_{n+1}} \\ &\times M_{n+3}^0(1, \dots, n+3) J_n^{(n+1)}(\{p\}_{n+1}), \end{aligned} \quad (2.3.2)$$

which denotes the integral of the real matrix element over the $(n+1)$ -parton phase space. The factor $\mathcal{N}_{\text{NLO}}^R$ encompasses all of the overall QCD and non-QCD factors relevant to this particular matrix element. It is related to the LO overall factor, \mathcal{N}_{LO} , by [68],

$$\mathcal{N}_{\text{NLO}}^R = \mathcal{N}_{\text{LO}} \left(\frac{\alpha_s(\mu_R^2)N}{2\pi} \right) \frac{\bar{C}(\epsilon)}{C(\epsilon)}. \quad (2.3.3)$$

The precise set of permutations considered depends on the particular colour level being considered. The coefficient $\mathcal{C}(N, n_f)$ is an overall colour factor relative to the leading colour contribution which is formed of powers of N and n_f . For matrix elements that distinguish between flavours of quarks on the basis of their electromagnetic charge, for

example matrix elements with intermediate Z/γ^* propagators, the colour factor has to be decomposed into ‘up-type’ n_{up} and ‘down-type’ n_{down} quark flavours. The full matrix element is then recovered once all colour contributions have been summed over.

For a particular colour level contribution to the full matrix element, the subtraction term is given by [68],

$$\begin{aligned} d\hat{\sigma}_{ij}^{S,\text{NLO}} &= \mathcal{N}_{\text{NLO}}^R \mathcal{C}(N, n_f) \sum_{\text{perms}} \sum_j d\Phi_{n+1}(p_3, \dots, p_{n+3}; p_1, p_2) \frac{1}{S_{n+1}} \\ &\times X_3^0(\dots, j, \dots) M_{n+2}^0(\dots, j, \dots) J_n^{(n)}(\{\tilde{p}_n\}), \end{aligned} \quad (2.3.4)$$

where the summation over j includes all final state partons.

The structure of the subtraction term is independent of labelling of the final state momentum. If all of the gluons are in the final state (for example $q\bar{q} \rightarrow$ gluons) then there is only one subtraction for the leading colour contribution. The full leading colour contribution can be recovered by fully permuting the momentum final state gluons. In less symmetric initial state configurations (for example $qg \rightarrow q +$ gluons) then a unique subtraction term needs to be constructed for each position the initial state gluon can be in the colour ordering. For NNLO subtraction terms typically colour orderings are grouped into topologies based on the symmetries of the antenna used in the process.

2.3.2 Virtual subtraction

Given the constraints of the KLN theorem, it is clear that the antenna structures introduced in the real phase space integral should be subtracted from the virtual contribution. To that end, the virtual matrix element must be colour decomposed in a similar fashion to the real $(n+1)$ -parton tree level matrix element. The integral to be considered is,

$$\int_{d\sigma_n} \left[d\hat{\sigma}_{ij}^V - d\hat{\sigma}_{ij}^{T,\text{NLO}} \right], \quad (2.3.5)$$

where the mass factorisation term has been absorbed into the definition of $d\hat{\sigma}_{ij}^{T,\text{NLO}}$, as demonstrated in Eq. (1.7.20). The integral of the virtual subtraction term is given by,

$$\begin{aligned} d\hat{\sigma}_{ij}^V &= \mathcal{N}_{\text{NLO}}^V \sum_{\mathcal{C}} \mathcal{C}(N, n_f) \sum_{\text{perms}} d\Phi_n(p_3, \dots, p_{n+2}; p_1, p_2) \frac{1}{S_n} \\ &\times M_{n+2}^1(1, \dots, n+2) J_n^{(n)}(\{p\}_n), \end{aligned} \quad (2.3.6)$$

| real correction colour string | Integrated dipole $\mathbf{J}_2^{(1)}$ | reduced colour string |
|----------------------------------|---|--------------------------|
| $(i_q, j_g, k_{\bar{q}})$ | $\mathbf{J}_{2,QQ}^{1,FF}(s_{IK}) = +\mathcal{A}_3^0(s_{IK})$ | $(I_q, K_{\bar{q}})$ |
| (i_q, j_g, k_g) | $\mathbf{J}_{2,QG}^{1,FF}(s_{IK}) = +\frac{1}{2} \mathcal{D}_3^0(s_{IK})$ | (I_q, K_g) |
| $(i_q, j_{\bar{q}'}; k_{q'})$ | $\hat{\mathbf{J}}_{2,QG}^{1,FF}(s_{IK}) = +\frac{1}{2} \mathcal{E}_3^0(s_{IK})$ | (I_q, K_g) |
| (i_g, j_g, k_g) | $\mathbf{J}_{2,GG}^{1,FF}(s_{IK}) = +\frac{1}{3} \mathcal{F}_3^0(s_{IK})$ | (I_g, K_g) |
| $(i_g, j_{\bar{q}}; k_q)$ | $\hat{\mathbf{J}}_{2,GG}^{1,FF}(s_{IK}) = +\mathcal{G}_3^0(s_{IK})$ | (I_g, K_g) |

Table 2.3: The definition of the final-final $\mathbf{J}_2^{(1)}$ dipoles in terms of integrated antenna and the corresponding singularity structure they subtract.

where as discussed in section 1.6.1, there are Catani operators consisting of explicit poles in ϵ between partons that are neighbouring in colour ordering. $\mathcal{N}_{\text{NLO}}^V$ is given by [68],

$$\mathcal{N}_{\text{NLO}}^V = \mathcal{N}_{\text{NLO}}^R \mathcal{C}(\epsilon). \quad (2.3.7)$$

In a similar approach to the real subtraction term, the subtraction term for a given colour ordering is,

$$\begin{aligned} d\hat{\sigma}_{ij}^{T,\text{NLO}} &= -\mathcal{N}_{\text{NLO}}^V \mathcal{C}(N, n_f) \int \frac{dx_1}{x_1} \frac{dx_2}{x_2} d\Phi_n(p_3, \dots, p_{n+2}; x_1 p_1, x_2 p_2) \frac{1}{S_n} \\ &\quad \times \sum_{I,K} \mathbf{J}_2^{(1)}(s_{IK}) M_{n+2}^0(\dots, I, K, \dots) J_n^{(n)}(\{p_n\}), \end{aligned} \quad (2.3.8)$$

where the $\mathbf{J}_2^{(1)}(s_{IK})$ functions are integrated dipole functions. The summation over I and K runs along the colour ordered string. The $\mathbf{J}_2^{(1)}(s_{IK})$ functions consist of an integrated antenna function $\mathcal{X}_3^0(s_{IK})$ and/or a mass factorisation splitting kernel $\Gamma_{ij,kl}(x_1, x_2)$. The precise composition of a given $\mathbf{J}_2^{(1)}(s_{IK})$ antenna function depends on the flavour content, colour level and initial state configuration of the antenna. They are given explicitly in Tabs. 2.3, 2.4 and 2.5 for final-final, initial-final and initial-initial integrated antenna functions respectively. For flavour preserving antenna functions, each $\mathbf{J}_2^{(1)}(s_{IK})$ is constructed to cancel the corresponding Catani operator at the same scale. This, combined with the mass factorisation contribution, renders the virtual subtraction term free of explicit ϵ poles.

The flavour changing $\mathbf{J}_2^{(1)}(s_{IK})$ integrated dipole functions are all finite because the pole structure of the integrated antenna function cancels entirely against the cor-

| real correction colour string | Integrated dipole $\mathbf{J}_2^{(1)}$ | | reduced colour string |
|---|--|---|----------------------------|
| Flavour preserving antenna | | | |
| $(\hat{1}_q, i_g, j_{\bar{q}})$ | $\mathbf{J}_{2,QQ}^{1,IF}(s_{1J})$ | $= +\mathcal{A}_{3,q}^0(s_{1J}) - \Gamma_{qq}^{(1)}(x_1)\delta_2$ | $(\hat{1}_q, J_{\bar{q}})$ |
| $(\hat{1}_q, i_g, j_g)$ | $\mathbf{J}_{2,QG}^{1,IF}(s_{1J})$ | $= +\frac{1}{2}\mathcal{D}_{3,q}^0(s_{1J}) - \Gamma_{qq}^{(1)}(x_1)\delta_2$ | $(\hat{1}_q, J_g)$ |
| $(\hat{1}_q, i_{\bar{q}'}; j_{q'})$ | $\hat{\mathbf{J}}_{2,QG}^{1,IF}(s_{1J})$ | $= +\frac{1}{2}\mathcal{E}_{3,q}^0(s_{1J})$ | $(\hat{1}_q, J_g)$ |
| $(i_q, j_g, \hat{1}_g)$ | $\mathbf{J}_{2,GQ}^{1,IF}(s_{1J})$ | $= +\mathcal{D}_{3,g\rightarrow g}^0(s_{1J}) - \frac{1}{2}\Gamma_{gg}^{(1)}(x_1)\delta_2$ | $(J_q, \hat{1}_g)$ |
| None | $\hat{\mathbf{J}}_{2,GQ}^{1,IF}(s_{1J})$ | $= -\frac{1}{2}\Gamma_{gg,F}^{(1)}(x_1)\delta_2$ | $(I_q, \hat{1}_g)$ |
| $(\hat{1}_g, i_g, j_g)$ | $\mathbf{J}_{2,GG}^{1,IF}(s_{1J})$ | $= +\frac{1}{2}\mathcal{F}_{3,g}^0(s_{1J}) - \frac{1}{2}\Gamma_{gg}^{(1)}(x_1)\delta_2$ | $(\hat{1}_g, J_g)$ |
| $(\hat{1}_g, i_{\bar{q}}; j_{\bar{q}})$ | $\hat{\mathbf{J}}_{2,GG}^{1,IF}(s_{1J})$ | $= +\frac{1}{2}\mathcal{G}_{3,g}^0(s_{1J}) - \frac{1}{2}\Gamma_{gg,F}^{(1)}(x_1)\delta_2$ | $(\hat{1}_g, J_g)$ |
| Flavour changing antenna | | | |
| $(i_q, \hat{1}_g, j_{\bar{q}})$ | $\mathbf{J}_{2,QQ,g\rightarrow q}^{1,IF}(s_{1J})$ | $= -\frac{1}{2}\mathcal{A}_{3,g\rightarrow q}^0(s_{1J}) - S_{g\rightarrow q}\Gamma_{qq}^{(1)}(x_1)\delta_2$ | $(\hat{1}_q, J_{\bar{q}})$ |
| $(i_q, \hat{1}_g, j_g)$ | $\mathbf{J}_{2,QG,g\rightarrow q}^{1,IF}(s_{1J})$ | $= -\mathcal{D}_{3,g\rightarrow q}^0(s_{1J}) - S_{g\rightarrow q}\Gamma_{qq}^{(1)}(x_1)\delta_2$ | $(\hat{1}_q, J_g)$ |
| $(i_{\bar{q}}, \hat{1}_{\bar{q}'}; j_{q'})$ | $\mathbf{J}_{2,GQ,q'\rightarrow g}^{1,IF}(s_{1J})$ | $= -\mathcal{E}_{3,q'\rightarrow g}^0(s_{1J}) - S_{q\rightarrow g}\Gamma_{qq}^{(1)}(x_1)\delta_2$ | $(J_{\bar{q}}, \hat{1}_g)$ |
| $(i_g, j_{\bar{q}}; \hat{1}_q)$ | $\mathbf{J}_{2,GG,q'\rightarrow g}^{1,IF}(s_{1J})$ | $= -\mathcal{G}_{3,q'\rightarrow g}^0(s_{1J}) - S_{q\rightarrow g}\Gamma_{qq}^{(1)}(x_1)\delta_2$ | $(J_g, \hat{1}_g)$ |

Table 2.4: The definition of the initial-final $\mathbf{J}_2^{(1)}$ dipoles in terms of integrated antenna and the corresponding singularity structure they subtract. For brevity $\delta_2 = \delta(1-x_2)$.

responding mass factorisation term. The factors $S_{q\rightarrow g}$ and $S_{g\rightarrow q}$ are defined as,

$$S_{g\rightarrow q} = 1 - \epsilon, \quad (2.3.9)$$

$$S_{q\rightarrow g} = \frac{1}{1 - \epsilon}, \quad (2.3.10)$$

and come from the relative factors of $1 - \epsilon$ from the initial-final flavour changing spin-averaging factors. Given that the flavour changing $\mathbf{J}_2^{(1)}$ functions are completely finite, it is important that a consistent overall factor is used between the unintegrated and integrated antenna functions. This is particularly important to check given that the final state averaging factors are typically different between the real and virtual contributions (S_{n+1} and S_n respectively).

For some of the n_f contributions the corresponding $\mathbf{J}_2^{(1)}$ dipole function does not have an antenna contribution, for example $\hat{\mathbf{J}}_{2,GQ}^{1,IF}(s_{1J})$ from Tab. 2.4. These contributions occur when there is no corresponding real correction at that colour level with that

| real correction colour string | Integrated dipole $\mathbf{J}_2^{(1)}$ | | reduced colour string |
|--|---|--|----------------------------------|
| Flavour preserving antenna | | | |
| $(\hat{1}_q, i_g, \hat{2}_{\bar{q}})$ | $\mathbf{J}_{2,QQ}^{1,II}(s_{12})$ | $= + \mathcal{A}_{3,qq}^0(s_{12}) - \Gamma_{qq}^{(1)}(x_1)\delta_2 - \Gamma_{qq}^{(1)}(x_2)\delta_1$ | $(\hat{1}_q, \hat{2}_{\bar{q}})$ |
| $(\hat{1}_q, i_g, \hat{2}_g)$ | $\mathbf{J}_{2,QG}^{1,II}(s_{12})$ | $= + \mathcal{D}_{3,gg}^0(s_{12}) - \Gamma_{qq}^{(1)}(x_1)\delta_2 - \frac{1}{2} \Gamma_{gg}^{(1)}(x_2)\delta_1$ | $(\hat{1}_q, \hat{2}_g)$ |
| None | $\hat{\mathbf{J}}_{2,QG}^{1,II}(s_{12})$ | $= - \frac{1}{2} \Gamma_{gg,F}^{(1)}(x_2)\delta_1$ | $(\hat{1}_q, \hat{2}_g)$ |
| $(\hat{1}_g, i_g, \hat{2}_g)$ | $\mathbf{J}_{2,GG}^{1,II}(s_{12})$ | $= + \mathcal{F}_{3,gg}^0(s_{12}) - \frac{1}{2} \Gamma_{gg}^{(1)}(x_1)\delta_2 - \frac{1}{2} \Gamma_{gg}^{(1)}(x_2)\delta_1$ | $(\hat{1}_g, \hat{2}_g)$ |
| None | $\hat{\mathbf{J}}_{2,GG}^{1,II}(s_{12})$ | $= - \frac{1}{2} \Gamma_{gg,F}^{(1)}(x_1)\delta_2 - \frac{1}{2} \Gamma_{gg,F}^{(1)}(x_2)\delta_1$ | $(\hat{1}_g, \hat{2}_g)$ |
| Flavour changing antenna | | | |
| $(\hat{1}_q, \hat{2}_g, i_{\bar{q}})$ | $\mathbf{J}_{2,QQ,qq \rightarrow qq}^{1,II}(s_{12})$ | $= - \mathcal{A}_{3,gg \rightarrow qq}^0(s_{12}) - S_{g \rightarrow q} \Gamma_{qq}^{(1)}(x_2)\delta_1$ | $(\hat{1}_q, \hat{2}_{\bar{q}})$ |
| $(i_q, \hat{1}_g, \hat{2}_g)$ | $\mathbf{J}_{2,QG,gg \rightarrow qq}^{1,II}(s_{12})$ | $= - \mathcal{D}_{3,gg \rightarrow qq}^0(s_{12}) - S_{g \rightarrow q} \Gamma_{qq}^{(1)}(x_1)\delta_2$ | $(\hat{1}_q, \hat{2}_g)$ |
| $(\hat{1}_q, \hat{2}_{\bar{q}}; i_{q'})$ | $\mathbf{J}_{2,QG,qq' \rightarrow qq}^{1,II}(s_{12})$ | $= - \mathcal{E}_{3,qq' \rightarrow qq}^0(s_{12}) - S_{q \rightarrow g} \Gamma_{qq}^{(1)}(x_2)\delta_1$ | $(\hat{1}_q, \hat{2}_g)$ |
| $(\hat{1}_g, \hat{2}_{\bar{q}}; i_q)$ | $\mathbf{J}_{2,GG,qq \rightarrow gg}^{1,II}(s_{12})$ | $= - \mathcal{G}_{3,qq \rightarrow gg}^0(s_{12}) - S_{q \rightarrow g} \Gamma_{qq}^{(1)}(x_2)\delta_1$ | $(\hat{1}_g, \hat{2}_g)$ |

Table 2.5: The definition of the initial-initial $\mathbf{J}_2^{(1)}$ dipoles in terms of integrated antenna and the corresponding singularity structure they subtract. For brevity $\delta_i = \delta(1 - x_i)$ for $i = 1, 2$.

specific initial state configuration. In this case the mass factorisation term will cancel the explicit singularity term from the one-loop matrix element contribution proportional to n_f (typically coming from the renormalisation counter term of the one-loop matrix element).

Another point to note is that there are no $n_f q\bar{q}$ dipole functions (e.g. $\hat{\mathbf{J}}_{2,QQ}^{1,II}(s_{12})$). This is because the only n_f contribution to the tree level splitting functions comes from $P_{gg}^{(0)}$ [30], hence only antennae with initial state gluons have an n_f contribution. At NNLO this is no longer true and the mass factorisation has a significantly more complicated colour structure.

By considering all colour levels for a given process, the antenna formalism regulates the implicit phase space divergences of the real integration, the explicit divergences

coming from the virtual correction and the mass factorisation counter terms whilst maintaining the consistency condition of Eq. (1.7.20).

2.4 Antenna subtraction at NNLO

As discussed at the beginning of this chapter, the antenna subtraction formalism can be used to calculate corrections at NNLO in perturbative QCD for processes with final state and initial state QCD radiation. In this section the general structure of the subtraction terms will be considered. In general, the structure of subtraction terms follows the same pattern for an arbitrary parton multiplicity. In some special cases the NNLO antenna functions contain additional limits or colour connections that require additional structures to regulate the spurious singularities. These structures will be discussed in section 2.5.

2.4.1 Double real subtraction

As demonstrated in the section 1.7.2, the double real integral contribution to an NNLO calculation is of the form,

$$\int_{d\sigma_{n+2}} \left[d\hat{\sigma}_{ij}^{RR} - d\hat{\sigma}_{ij}^S \right], \quad (2.4.1)$$

where $d\hat{\sigma}_{ij}^{RR}$ is a straightforward extension to the real parton integration definition,

$$\begin{aligned} d\hat{\sigma}_{ij}^{RR} &= \mathcal{N}_{\text{NNLO}}^{RR} \sum_{\mathcal{C}} \mathcal{C}(N, n_f) \sum_{\text{perms}} d\Phi_{n+2}(p_3, \dots, p_{n+4}; p_1, p_2) \frac{1}{S_{n+2}} \\ &\times M_{n+4}^0(1, \dots, n+4) J_n^{(n+2)}(\{p\}_{n+2}), \end{aligned} \quad (2.4.2)$$

where now the jet algorithm must form at least n jets from an $(n+2)$ -parton final state. The overall factor, $\mathcal{N}_{\text{NNLO}}^{RR}$, is defined as [68],

$$\mathcal{N}_{\text{NNLO}}^{RR} = \mathcal{N}_{\text{LO}} \left(\frac{\alpha_s(\mu_R^2) N}{2\pi} \right)^2 \frac{\bar{C}(\epsilon)^2}{C(\epsilon)^2}. \quad (2.4.3)$$

The double real subtraction term, $d\hat{\sigma}_{ij}^S$, must regulate all single unresolved and double unresolved singularity structures in $d\hat{\sigma}_{ij}^{RR}$. To illustrate how $d\hat{\sigma}_{ij}^S$ performs this task, the subtraction term will be broken down into different sections,

$$d\hat{\sigma}_{ij}^S = d\hat{\sigma}_{ij}^{S,a} + d\hat{\sigma}_{ij}^{S,b_1} + d\hat{\sigma}_{ij}^{S,b_2} + d\hat{\sigma}_{ij}^{S,c} + d\hat{\sigma}_{ij}^{S,d}, \quad (2.4.4)$$

each performing a different role in isolating the singularities of the process.

Single unresolved emissions $d\hat{\sigma}_{ij}^{S,a}$

The first section of the double real subtraction term, $d\hat{\sigma}_{ij}^{S,a}$, follows the same structure as the real subtraction term at NLO, only with an additional parton in the final state. This is to regulate the single unresolved limits of the double real matrix element. Explicitly this is given by,

$$d\hat{\sigma}_{ij}^{S,a} = \mathcal{N}_{\text{NNLO}}^{RR} \mathcal{C}(N, n_f) \sum_{\text{perms}} \sum_j d\Phi_{n+2}(p_3, \dots, p_{n+4}; p_1, p_2) \frac{1}{S_{n+2}} \\ \times X_3^0(\dots, j, \dots) M_{n+3}^0(\dots, j, \dots) J_n^{(n+1)}(\{\tilde{p}\}_{n+1}), \quad (2.4.5)$$

for a specific colour level. The jet algorithm is applied to the $(n+1)$ -parton factorised phase space.

Double unresolved colour connected emissions $d\hat{\sigma}_{ij}^{S,b_1}$

The next section of the subtraction term regulates colour connected double unresolved emissions. As introduced earlier, double unresolved emissions can be regulated using $X_4^0(i, j, k, l)$ antenna functions. Explicitly the structure introduced is of the form,

$$d\hat{\sigma}_{ij}^{S,b_1} = \mathcal{N}_{\text{NNLO}}^{RR} \mathcal{C}(N, n_f) \sum_{\text{perms}} \sum_{j,k} d\Phi_{n+2}(p_3, \dots, p_{n+4}; p_1, p_2) \frac{1}{S_{n+2}} \\ \times X_4^0(\dots, j, k, \dots) M_{n+2}^0(\dots, j, k, \dots) J_n^{(n)}(\{\tilde{p}\}_n). \quad (2.4.6)$$

Double unresolved overlapping correction $d\hat{\sigma}_{ij}^{S,b_2}$

Two issues have been introduced by including the single unresolved subtraction term and the colour connected double unresolved emissions,

- in double unresolved colour connected limits, the mapped matrix elements $M_{n+3}^0(\dots, j, \dots)$ are divergent. The single unresolved $X_3^0(\dots, j, \dots)$ antenna function only regulates one of the two divergent limits in this case.
- The $X_4^0(\dots, j, k, \dots)$ antenna functions contain spurious singularities in single unresolved limits.

A new series of terms must be implemented to regulate the spurious singularities introduced by these two blocks of terms. This can be performed by introducing terms with two $X_3^0(\dots, j, \dots)$ antenna functions [68],

$$\begin{aligned} d\hat{\sigma}_{ij}^{S,b_2} &= -\mathcal{N}_{\text{NNLO}}^{RR} \mathcal{C}(N, n_f) \sum_{\text{perms}} \sum_j d\Phi_{n+2}(p_3, \dots, p_{n+4}; p_1, p_2) \frac{1}{S_{n+2}} \\ &\times X_3^0(i, j, k) X_3^0(I, K, l) M_{n+2}^0(\dots, I', L, \dots) J_n^{(n)}(\{\tilde{p}\}_n), \end{aligned} \quad (2.4.7)$$

where the summation over j now runs along the colour connected partons within the $X_4^0(i, j, k, l)$ antenna instead of the full matrix element. The momentum set $\{\tilde{p}\}_n$ denotes the momentum configuration following both mappings.

Almost colour connected unresolved emissions $d\hat{\sigma}_{ij}^{S,c}$

In processes involving five or more partons, it is possible to generate almost colour connected double unresolved emissions, where one hard radiator separates two unresolved emissions in a colour ordering. In these limits, the terms introduced thus far exhibit two problems,

- the terms in $d\hat{\sigma}_{ij}^{S,a}$ double count the correct singularity structure. Consider the colour structure (i, j, k, l, m) , where partons j and l go unresolved simultaneously. In $d\hat{\sigma}_{ij}^{S,a}$ there will be a term introduced for when j is in a single unresolved limit and a second line for when l is unresolved. Both of these terms correctly subtract the almost colour connected singularity structure and thus the limit is double counted.
- A specific classification of X_4^0 antenna, denoted by \widetilde{X}_4^0 , exhibits an oversubtraction problem in almost colour connected limits. The terms in $d\hat{\sigma}_{ij}^{S,b_2}$ typically introduce a spurious limit in this configuration for \widetilde{X}_4^0 antennae. An example being the $D_4^0(\hat{1}_q, j_g, k_g, l_g)$ antenna, where using precisely the same argument as above, the $j_g || k_g, \hat{1}_q || l_g$ double collinear limit is double counted by the terms in $d\hat{\sigma}_{ij}^{S,b_2}$. Indeed, this issue arises because one of the underlying colour structures of the $D_4^0(\hat{1}_q, j_g, k_g, l_g)$ antenna is³ $(\hat{1}_q, j_g, k_g, l_g, \hat{1}_{\bar{q}})$. When we remove all of the

³For the sake of this discussion the additional colour connection as a result of the quark having the colour connections of a gluino will be ignored. This can be regulated using a different approach

| Mapping | X_4^0 | \widetilde{X}_4^0 |
|-----------------|---|--|
| final-final | $A_4^0, B_4^0, C_4^0, D_{4,a}^0, E_{4,a}^0, \widetilde{E}_4^0, F_{4,a}^0, G_{4,a}^0, G_{4,b}^0, \widetilde{G}_{4,a}^0, H_4^0$ | $\widetilde{A}_4^0, D_{4,c}^0, E_{4,b}^0, F_{4,b}^0, G_{4,c}^0$ |
| initial-final | $A_4^0, B_4^0, C_4^0, G_4^0, \widetilde{G}_4^0, H_4^0$ | $\widetilde{A}_4^0, D_4^0, E_4^0, F_4^0$ |
| initial-initial | $A_4^0, B_4^0, C_4^0, D_{4,\text{adj}}^0, F_{4,\text{adj}}^0, G_4^0, \widetilde{G}_4^0, H_4^0$ | $\widetilde{A}_4^0, D_{4,\text{n.adj}}^0, E_4^0, F_{4,\text{n.adj}}^0$ |

Table 2.6: The classification of double unresolved antenna without almost colour connected limits (X_4^0) and those with almost colour connected limits (\widetilde{X}_4^0) [68].

single unresolved limits of this colour ordering using the terms of $d\hat{\sigma}_{ij}^{S,b_2}$, the limits that connect the two hard radiators will be double counted. A full list of the \widetilde{X}_4^0 antenna are given in Tab. 2.6.

Extra terms must therefore be introduced to regulate the almost colour connected limits. This can be achieved by following the prescription established in Ref. [68]. Firstly, consider a matrix element with underlying parton colour ordering $(\dots, a, i, j, k, l, b, \dots)$, where within this process the antenna $\widetilde{X}_4^0(i, j, k, l)$ has been introduced. As usual, partons i, l are hard emitters and j, k are the unresolved emissions. The underlying colour ordering is therefore $(\dots, a, i, l, b, \dots)$, where the unresolved emissions have been removed. Then consider the scenario where parton k is emitted between i and l , leading to a colour ordering of $(\dots, a, i, k, l, b, \dots)$. There are two possible almost colour connected orderings between parton j and parton k ,

$$(\dots, a, \mathbf{j}, i, \mathbf{k}, l, b, \dots), \quad (2.4.8)$$

$$(\dots, a, i, \mathbf{k}, l, \mathbf{j}, b, \dots), \quad (2.4.9)$$

(where the unresolved partons are in bold for illustrative purposes). The $\widetilde{X}_4^0(i, j, k, l)$ antenna introduces an unphysical almost colour connected ordering of,

$$(i, \mathbf{j}, \mathbf{k}, l, i), \quad (2.4.10)$$

hence the general structure to subtract these limits for the final-final configuration is given by [68],

$$+ \frac{1}{2} X_3^0(i, j, l) X_3^0(\widetilde{(ij)}, k, \widetilde{(jl)}) M_n^0(\dots, a, \widetilde{((ij)k)}, \widetilde{(k(jl))}, b, \dots)$$

which will be discussed in section 2.5.

$$\begin{aligned}
& -\frac{1}{2}X_3^0(a, j, i)X_3^0(\widetilde{(ij)}, k, l)M_n^0(\dots, \widetilde{(aj)}, \widetilde{((ij)k)}, \widetilde{(kl)}, b, \dots) \\
& -\frac{1}{2}X_3^0(b, j, l)X_3^0(i, k, \widetilde{(jl)})M_n^0(\dots, a, \widetilde{(ik)}, \widetilde{(k(lj))}, \widetilde{(jb)}, \dots) + (j \leftrightarrow k), \quad (2.4.11)
\end{aligned}$$

where the first line regulates the double counting of the limits from $d\hat{\sigma}_{ij}^{S, b_2}$ and the final two lines regulate the double counting from $d\hat{\sigma}_{ij}^{S, a}$. To complete this block of antenna partons j and k must be swapped. This structure is extended trivially to processes involving initial state partons.

The introduction of these terms leads to yet another oversubtraction of the form of wide angle soft terms [68, 100, 101]. In all of the lines of Eq. (2.4.11) there can be divergences when either j or k are soft. In these single soft limits there are no antenna structures that cancel these spurious limits, forcing us to introduce explicit Eikonal functions to cancel these limits.

Consider the j soft limit for the terms in Eq. (2.4.11). All six terms are divergent in this limit, where three of the divergent Eikonal factors have not undergone a phase space mapping and three have. In soft limits all of the mapping functions factorise in precisely the same way. This means that unlike the antenna in Eq. (2.4.11), a common first mapping can be chosen for all of the eikonal factors [68],

$$Y \cdot X_3^0(\widetilde{(ij)}, k, \widetilde{(jl)})M_n^0(\dots, a, \widetilde{((ij)k)}, \widetilde{(k(jl))}, b, \dots), \quad (2.4.12)$$

where Y is a function that includes all of the necessary wide angle soft terms. It follows a similar format to Eq. (2.4.11) and is of the form,

$$\frac{1}{2}(S_{\alpha j \beta} - S_{A J B}), \quad (2.4.13)$$

where J is parton j after the second mapping, α and β are the hard radiators after the first mapping and A and B are the hard radiators following the second mapping. Including the soft terms allows us to construct the complete almost colour connected subtraction contribution,

$$\begin{aligned}
d\hat{\sigma}_{ij}^{S, c} &= \mathcal{N}_{\text{NNLO}}^{RR} \mathcal{C}(N, n_f) \sum_{j, k} d\Phi_{n+2}(p_3, \dots, p_{n+4}; p_1, p_2) \frac{1}{S_{n+2}} \left\{ \right. \\
& + \frac{1}{2}X_3^0(i, j, l)X_3^0(\widetilde{(ij)}, k, \widetilde{(jl)})M_n^0(\dots, a, \widetilde{((ij)k)}, \widetilde{(k(jl))}, b, \dots)J_n^{(n)}(\{\tilde{p}\}_n) \\
& - \frac{1}{2}X_3^0(a, j, i)X_3^0(\widetilde{(ij)}, k, l)M_n^0(\dots, \widetilde{(aj)}, \widetilde{((ij)k)}, \widetilde{(kl)}, b, \dots)J_n^{(n)}(\{\tilde{p}\}_n)
\end{aligned}$$

$$\begin{aligned}
& -\frac{1}{2}X_3^0(b, j, l)X_3^0(i, k, (\widetilde{jl}))M_n^0(\dots, a, (\widetilde{ik}), (\widetilde{k(lj)}), (\widetilde{jb}), \dots)J_n^{(n)}(\{\widetilde{p}\}_n) \\
& -\frac{1}{2}\left[(S_{(\widetilde{ij})j(\widetilde{jl})} - S_{(\widetilde{(ij)k})j(\widetilde{k(jl)})}) - (S_{a_j(\widetilde{ij})} - S_{a_i(\widetilde{k(ij)})}) - (S_{b_j(\widetilde{jl})} - S_{b_j(\widetilde{k(jl)})})\right] \\
& \quad \times X_3^0((\widetilde{ij}), k, (\widetilde{jl}))M_n^0(\dots, a, (\widetilde{(ij)k}), (\widetilde{k(jl)}), b, \dots)J_n^{(n)}(\{\widetilde{p}\}_n) \\
& \quad \left. + (j \leftrightarrow k) \right\}, \tag{2.4.14}
\end{aligned}$$

where in this case the first mapping for the soft terms was chosen to be the final-final mapping function. The choice of mapping function is reliant on two factors,

- there must be sufficient partons in the final state such that the hard radiators remain hard following both mappings. This implies that the final-final mapping can only be used when there are a minimum of six partons in the process.
- The Eikonal function must have been integrated analytically for the choice of phase space mapping. Only the initial-final [95] and final-final [67] configurations have been integrated analytically, meaning that the initial-initial Eikonal function can never be used. Nevertheless, the final-final and initial-final soft functions are sufficient for all parton multiplicities.

Colour disconnected unresolved emissions $d\hat{\sigma}_{ij}^{S,d}$

For processes with six or more partons, it is possible to generate completely colour disconnected emissions, where there are at least two hard radiators in the colour ordering between the two unresolved emissions. Again the terms in $d\hat{\sigma}_{ij}^{S,a}$ double count the limits in much the same way as the almost colour connected limits. This can be solved simply by introducing terms of the form,

$$\begin{aligned}
d\hat{\sigma}_{ij}^{S,d} &= -\mathcal{N}_{\text{NNLO}}^{RR} \mathcal{C}(N, n_f) \sum_{\text{perms}} \sum_j d\Phi_{n+2}(p_3, \dots, p_{n+4}; p_1, p_2) \frac{1}{S_{n+2}} \\
& \quad \times X_3^0(i, j, k)X_3^0(l, m, n)M_{n+2}^0(\dots, I, K, \dots, L, M, \dots)J_n^{(n)}(\{\widetilde{p}\}_n), \tag{2.4.15}
\end{aligned}$$

where given the colour disconnected nature of the limit the choice of primary antenna is irrelevant.

2.4.2 Real virtual subtraction

The real virtual contribution to the cross section is of the form,

$$\int_{d\sigma_{n+1}} \left[d\hat{\sigma}_{ij}^{RV} - d\hat{\sigma}_{ij}^T \right], \quad (2.4.16)$$

where $d\hat{\sigma}_{ij}^{RV}$ contains both explicit divergences coming from the loop integral in d -dimensions and implicit phase space singularities coming from an unresolved emission in the final state. It is defined as a simple extension of $d\hat{\sigma}_{ij}^V$ with one additional parton in the final state,

$$\begin{aligned} d\hat{\sigma}_{ij}^{RV} &= \mathcal{N}_{\text{NNLO}}^{RV} \sum_{\mathcal{C}} \mathcal{C}(N, n_f) \sum_{\text{perms}} d\Phi_{n+1}(p_3, \dots, p_{n+3}; p_1, p_2) \frac{1}{S_{n+1}} \\ &\times M_{n+3}^1(1, \dots, n+3) J_n^{(n+1)}(\{p\}_{n+1}), \end{aligned} \quad (2.4.17)$$

where $\mathcal{N}_{\text{NNLO}}^{RV} = C(\epsilon) \mathcal{N}_{\text{NNLO}}^{RR}$. As with the double real subtraction term, it is convenient to decompose $d\hat{\sigma}_{ij}^T$ into individual components,

$$d\hat{\sigma}_{ij}^T = d\hat{\sigma}_{ij}^{T,a} + d\hat{\sigma}_{ij}^{T,b_1} + d\hat{\sigma}_{ij}^{T,b_2} + d\hat{\sigma}_{ij}^{T,c}, \quad (2.4.18)$$

and discuss each component in turn.

Explicit pole cancellation, $d\hat{\sigma}_{ij}^{T,a}$

The first block of terms in the real virtual subtraction term, $d\hat{\sigma}_{ij}^{T,a}$, is designed to cancel the explicit poles arising from $d\hat{\sigma}_{ij}^{RV}$. This is directly analogous to the entirety of $d\hat{\sigma}_{ij}^{T,\text{NLO}}$ with one additional parton in the final state,

$$\begin{aligned} d\hat{\sigma}_{ij}^{T,a} &= -\mathcal{N}_{\text{NNLO}}^{RV} \mathcal{C}(N, n_f) \int \frac{dx_1}{x_1} \frac{dx_2}{x_2} \frac{\bar{C}(\epsilon)}{C(\epsilon)} d\Phi_{n+1}(p_3, \dots, p_{n+3}; x_1 p_1, x_2 p_2) \frac{1}{S_{n+1}} \\ &\times \sum_{I,K} \mathbf{J}_2^{(1)}(s_{IK}) M_{n+3}^0(\dots, I, K, \dots) J_n^{(n+1)}(\{p\}_{n+1}). \end{aligned} \quad (2.4.19)$$

Tree \times loop implicit singularities, $d\hat{\sigma}_{ij}^{T,b_1}$

As discussed in section 1.6.4, a one-loop matrix element in a single unresolved configuration factorises into a sum of two divergent contributions. To regulate the full implicit singularity structure of the one-loop matrix element, the subtraction term will also mimic this structure. $d\hat{\sigma}_{ij}^{T,b_1}$ contains a tree level antenna function factorising onto

a reduced one-loop matrix element. Given that the reduced matrix element contains an explicit ϵ pole structure, $\mathbf{J}_2^{(1)}$ integrated dipoles must be introduced in the reduced momentum set to regulate these explicit pole structures. Combining these contributions implies that $d\hat{\sigma}_{ij}^{T,b_1}$ is given by [68],

$$\begin{aligned} d\hat{\sigma}_{ij}^{T,b_1} &= \mathcal{N}_{\text{NNLO}}^{RV} \mathcal{C}(N, n_f) \int \frac{dx_1}{x_1} \frac{dx_2}{x_2} d\Phi_{n+1}(p_3, \dots, p_{n+3}; x_1 p_1, x_2 p_2) \frac{1}{S_{n+1}} \\ &\times \sum_j X_3^0(i, j, k) \left\{ \delta(1-x_1)\delta(1-x_2) M_{n+2}^1(\dots, I, K, \dots) \right. \\ &\quad \left. + \sum_{I, K} \mathbf{J}_2^{(1)}(s_{IK}) M_{n+2}^0(\dots, I, K, \dots) \right\} J_n^{(n)}(\{\tilde{p}\}_n), \end{aligned} \quad (2.4.20)$$

where the sum over I, K runs along colour connected partons in the reduced momentum set, implying that $d\hat{\sigma}_{ij}^{T,b_1}$ is free of explicit ϵ poles⁴.

loop \times tree implicit singularities, $d\hat{\sigma}_{ij}^{T,b_2}$

The regulation of the implicit singularity structure of one-loop matrix element is completed by a one-loop antenna function factorising onto a tree level reduced matrix element. In this case the one-loop antenna function contains explicit ϵ poles that must be regulated by $\mathbf{J}_2^{(1)}$ integrated dipole functions. Explicitly this is given by [68],

$$\begin{aligned} d\hat{\sigma}_{ij}^{T,b_2} &= \mathcal{N}_{\text{NNLO}}^{RV} \mathcal{C}(N, n_f) \int \frac{dx_1}{x_1} \frac{dx_2}{x_2} \frac{\bar{C}(\epsilon)}{C(\epsilon)} d\Phi_{n+1}(p_3, \dots, p_{n+3}; x_1 p_1, x_2 p_2) \frac{1}{S_{n+1}} \\ &\times \sum_j \left[X_3^1(i, j, k) \delta(1-x_1)\delta(1-x_2) M_{n+2}^0(\dots, I, K, \dots) \right. \\ &\quad + \sum_{i, j} \mathbf{J}_2^{(1)}(s_{ij}) X_3^0(i, j, k) M_{n+2}^0(\dots, I, K, \dots) \\ &\quad \left. - M_X X_3^0(i, j, k) \mathbf{J}_2^{(1)}(s_{IK}) M_{n+2}^0(\dots, I, K, \dots) \right] J_n^{(n)}(\{\tilde{p}\}_n), \end{aligned} \quad (2.4.21)$$

where the sum over i, j indicates a sum over the colour connected partons in the $X_3^1(i, j, k)$ antenna function. The $\mathbf{J}_2^{(1)}(s_{ij})$ integrated dipoles are defined in the full momentum set whereas $\mathbf{J}_2^{(1)}(s_{IK})$ is defined in the mapped momentum set. M_X is a constant that is either 0, 1 or 2 depending on the $X_3^1(i, j, k)$ under consideration.

⁴There is a special case where $n = 0$ and the reduced matrix element contains only gluons. In this case an additional factor of 2 for the $\mathbf{J}_2^{(1)}(s_{IK})$ is needed to fully subtract the explicit poles.

The final line of Eq. (2.4.21) is required to regulate the explicit pole structure of the subtracted ‘tree \times loop’ structure from $X_3^1(i, j, k)$. The reason why the coefficient is dependent on the antenna in question is because in some colour configurations there is no reduced loop matrix element, hence $M_X = 0$ for those antenna. The case of $M_X = 2$ arises in antenna that factorise onto a pure gluon matrix element.

Given that the $X_3^1(i, j, k)$ antenna functions were derived from one-loop matrix elements, it is necessary to also introduce a renormalisation counter term. The $X_3^1(i, j, k)$ antenna are renormalised at the scale s_{ijk} , given that it is the only physically appropriate choice, whereas the subtraction term is renormalised at another scale μ_R^2 . To ensure that all of the terms are renormalised to the same scale, a redefinition of the $X_3^1(i, j, k)$ antenna functions is required of the form [68],

$$X_3^1(i, j, k) \rightarrow X_3^1(i, j, k) + \frac{\beta_0}{\epsilon} X_3^0(i, j, k) \left[\left(\frac{|s_{ijk}|}{\mu_R^2} \right)^{-\epsilon} - 1 \right]. \quad (2.4.22)$$

Following this redefinition, $d\hat{\sigma}_{ij}^{T,b2}$ is free of explicit ϵ poles.

Integrated almost colour connected subtraction term, $d\hat{\sigma}_{ij}^{T,c}$

For processes that do not have an almost colour connected contribution in the double real subtraction term, the above contributions are sufficient to regulate all of the implicit and explicit singularity structures of the one-loop matrix element. For these simple processes, the additional $\mathbf{J}_2^{(1)}$ integrated dipoles that were introduced in $d\hat{\sigma}_{ij}^{T,b1}$ and $d\hat{\sigma}_{ij}^{T,b2}$ are sufficient to regulate all of implicit poles of the reduced matrix elements in $d\hat{\sigma}_{ij}^{T,a}$. For processes with more partons however this is not sufficient and we need to integrate the terms from $d\hat{\sigma}_{ij}^{S,c}$ to regulate the remaining implicit singularities from $d\hat{\sigma}_{ij}^{T,a}$. Explicitly this is given by [68],

$$d\hat{\sigma}_{ij}^{T,c} = -\mathcal{N}_{\text{NNLO}}^{RV} \mathcal{C}(N, n_f) \sum_{j,k} d\Phi_{n+2}(p_3, \dots, p_{n+3}; p_1, p_2) \frac{1}{S_{n+1}} \left\{ \right.$$

$$\begin{aligned}
& + \frac{1}{2} \sum_j \left[+ (\mathcal{X}_3^0(s_{ik}) - \mathcal{X}_3^0(s_{IK})) \right. \\
& \quad - (\mathcal{X}_3^0(s_{ai}) - \mathcal{X}_3^0(s_{aI})) \\
& \quad - (\mathcal{X}_3^0(s_{bk}) - \mathcal{X}_3^0(s_{bK})) \\
& \quad - \left(+ (\mathcal{S}(s_{ik}, s_{ik}, 1) - \mathcal{S}(s_{IK}, s_{IK}, x_{IK,ik})) \right. \\
& \quad \quad - (\mathcal{S}(s_{ai}, s_{ik}, x_{ai,ik}) - \mathcal{S}(s_{aI}, s_{ik}, x_{aI,ik})) \\
& \quad \quad \left. \left. - (\mathcal{S}(s_{bk}, s_{ik}, x_{bk,ik}) - \mathcal{S}(s_{bK}, s_{ik}, x_{bK,ik})) \right) \delta(1-x_1)\delta(1-x_2) \right] \Big\} \\
& \times X_3^0(i, j, k) M_{n+2}^0(\dots, I, K, \dots) J_n^{(n)}(\{\tilde{p}\}_n), \tag{2.4.23}
\end{aligned}$$

where the integrated antennae that act on the reduced momentum set are introduced to regulate the poles of the antenna coming from $d\hat{\sigma}_{ij}^{S,c}$. These terms will be integrated out in the double virtual subtraction term. The functions $\mathcal{S}(s_{ai}, s_{ik}, x_{ai,ik})$ denote the integrated soft functions, either in the initial-final or final-final configuration, depending on the choice of mapping in the double real subtraction term. Similarly to $d\hat{\sigma}_{ij}^{T,b_1}$ and $d\hat{\sigma}_{ij}^{T,b_2}$, $d\hat{\sigma}_{ij}^{T,b_c}$ is free of explicit poles and regulates the remaining implicit singularities in the previous contributions to $d\hat{\sigma}_{ij}^T$.

2.4.3 Double virtual subtraction

The final contribution to the NNLO corrections to the cross section is the double virtual integral,

$$\int_{d\sigma_n} \left[d\hat{\sigma}_{ij}^{VV} - d\hat{\sigma}_{ij}^U \right], \tag{2.4.24}$$

where $d\hat{\sigma}_{ij}^{VV}$ is integral of the two-loop matrix element,

$$\begin{aligned}
d\hat{\sigma}_{ij}^{VV} &= \mathcal{N}_{\text{NNLO}}^{VV} \sum_{\mathcal{C}} \mathcal{C}(N, n_f) \sum_{\text{perms}} d\Phi_n(p_3, \dots, p_{n+2}; p_1, p_2) \frac{1}{S_n} \\
&\times M_{n+2}^2(1, \dots, n+2) J_n^{(n)}(\{p\}_n), \tag{2.4.25}
\end{aligned}$$

where $\mathcal{N}_{\text{NNLO}}^{VV} = C(\epsilon) \mathcal{N}_{\text{NNLO}}^{RV} = C(\epsilon)^2 \mathcal{N}_{\text{NNLO}}^{RR}$. As described in section 1.6.2, the explicit ϵ pole structure of the two-loop matrix element is described by a series of Catani operators. The construction of the double virtual subtraction term will therefore follow a similar structure as the Catani operators themselves. $d\hat{\sigma}_{ij}^U$ can be decomposed into

three sections,

$$d\hat{\sigma}_{ij}^U = d\hat{\sigma}_{ij}^{U,a} + d\hat{\sigma}_{ij}^{U,b} + d\hat{\sigma}_{ij}^{U,c}, \quad (2.4.26)$$

where each term regulates one or more lines of the two-loop Catani pole structure in Eq. (1.6.22).

One loop factorised term, $d\hat{\sigma}_{ij}^{U,a}$

$d\hat{\sigma}_{ij}^{U,a}$ is inherited from the ‘tree \times loop’ structure in $d\hat{\sigma}_{ij}^{T,b_1}$,

$$\begin{aligned} d\hat{\sigma}_{ij}^{U,a} &= -\mathcal{N}_{\text{NNLO}}^{VV} \mathcal{C}(N, n_f) \int \frac{dx_1}{x_1} \frac{dx_2}{x_2} d\Phi_n(p_3, \dots, p_{n+2}; p_1, p_2) \frac{1}{S_n} \\ &\times \sum_{i,j} \mathbf{J}_2^{(1)}(s_{ij}) \left[M_{n+2}^1(\dots, i, j, \dots) - \frac{\beta_0}{\epsilon} M_{n+2}^0(\dots, i, j, \dots) \right] J_n^{(n)}(\{p\}_n), \end{aligned} \quad (2.4.27)$$

which regulates the $1/\epsilon^4$ and $1/\epsilon^3$ poles in the first line of Eq. (1.6.22). The lower order pole structures of the first line of Eq. (1.6.22) are not described by $d\hat{\sigma}_{ij}^{U,a}$ however, given that the finite term in the $\mathbf{J}_2^{(1)}$ operators differs from the equivalent Catani operators.

To construct this contribution to the subtraction term, part of the redefinition of the $X_3^1(i, j, k)$ antenna has been absorbed into $d\hat{\sigma}_{ij}^{U,a}$. The contribution proportional to ‘-1’ in Eq. (2.4.22) has been extracted to generate the terms proportional to β_0/ϵ in $d\hat{\sigma}_{ij}^{U,a}$. The term proportional to the scale $(s_{ijk}/\mu_R^2)^{-\epsilon}$ will be absorbed into $d\hat{\sigma}_{ij}^{U,c}$.

Overlapping $\mathbf{J}_2^{(1)} \otimes \mathbf{J}_2^{(1)}$ term, $d\hat{\sigma}_{ij}^{U,b}$

The many terms that result in an overlapping singularity structure can be absorbed into terms proportional to $[\mathbf{J}_2^{(1)} \otimes \mathbf{J}_2^{(1)}](x_1, x_2)$, a convolution between two $\mathbf{J}_2^{(1)}$ operators. $d\hat{\sigma}_{ij}^{U,b}$ describes all of the contributions to the double virtual subtraction term that contain mixed scales and is given by,

$$\begin{aligned} d\hat{\sigma}_{ij}^{U,b} &= -\mathcal{N}_{\text{NNLO}}^{VV} \mathcal{C}(N, n_f) \int \frac{dx_1}{x_1} \frac{dx_2}{x_2} d\Phi_n(p_3, \dots, p_{n+2}; p_1, p_2) \frac{1}{S_n} \\ &\times \frac{1}{2} \left[\sum_{i,j} \mathbf{J}_2^{(1)}(s_{ij}) \otimes \sum_{k,l} \mathbf{J}_2^{(1)}(s_{kl}) \right] (x_1, x_2) M_{n+2}^0(\dots, i, j, \dots, k, l, \dots) J_n^{(n)}(\{p\}_n), \end{aligned} \quad (2.4.28)$$

which describes the $1/\epsilon^4$ and $1/\epsilon^3$ poles of the second line of Eq. (1.6.22). The combination of the lines $d\hat{\sigma}_{ij}^{U,a}$ and $d\hat{\sigma}_{ij}^{U,b}$ describes all of the $1/\epsilon^4$, $1/\epsilon^3$ and $1/\epsilon^2$ poles [68].

Tree level factorised term, $d\hat{\sigma}_{ij}^{U,c}$

$d\hat{\sigma}_{ij}^{U,c}$ is a collection of all of the remaining terms from $d\hat{\sigma}_{ij}^S$ and $d\hat{\sigma}_{ij}^T$, including the remaining NNLO mass factorisation contributions. In precisely the same format as $d\hat{\sigma}_{ij}^{T,\text{NLO}}$, the singularity structures can be described as strings of integrated dipole functions. In this case they are referred to as $\mathbf{J}_2^{(2)}$ operators. $d\hat{\sigma}_{ij}^{U,c}$ is given by,

$$d\hat{\sigma}_{ij}^{U,c} = -\mathcal{N}_{\text{NLO}}^{VV}\mathcal{C}(N, n_f) \int \frac{dx_1}{x_1} \frac{dx_2}{x_2} d\Phi_n(p_3, \dots, p_{n+2}; x_1 p_1, x_2 p_2) \frac{1}{S_n} \\ \times \sum_{i,j} \mathbf{J}_2^{(2)}(s_{ij}) M_{n+2}^0(\dots, i, j, \dots) J_n^{(n)}(\{p_n\}). \quad (2.4.29)$$

The $\mathbf{J}_2^{(2)}$ integrated dipole functions have many more components than their $\mathbf{J}_2^{(1)}$ counterparts. The most general $\mathbf{J}_2^{(2)}$ operator is of the form [68],

$$\mathbf{J}_2^{(2)}(s_{ij}) = c_1 \mathcal{X}_4^0(s_{ij}) + c_2 \tilde{\mathcal{X}}_4^0(s_{ij}) + c_3 \mathcal{X}_3^1(s_{ij}) \\ + c_4 \frac{\beta_0}{\epsilon} \left(\frac{|s_{ij}|}{\mu^2} \right)^{-\epsilon} \mathcal{X}_3^0(s_{ij}) + c_5 \mathcal{X}_3^0(s_{ij}) \otimes \mathcal{X}_3^0(s_{ij}) \\ - m_1 \bar{\Gamma}_{ij;kl}^{(2)}(x_1) \delta(1-x_2) - m_2 \bar{\Gamma}_{ij;kl}^{(2)}(x_2) \delta(1-x_1), \quad (2.4.30)$$

where the set of coefficients c_i are dependent on the integrated dipole function in question. The function $\bar{\Gamma}_{ij;kl}^{(2)}(x)$ describes all of the remaining mass factorisation contributions which were not included in $d\hat{\sigma}_{ij}^{U,a}$ and $d\hat{\sigma}_{ij}^{U,b}$. A detailed discussion of the decomposition of NNLO mass factorisation terms into the antenna subtraction structure can be found in Ref. [68]. For initial-initial $\mathbf{J}_2^{(2)}$ dipoles $m_1 = m_2 = 1$, for initial-final $m_1 = 1, m_2 = 0$ and final-final dipoles $m_1 = m_2 = 0$.

Once all of the remaining terms have been collected into $\mathbf{J}_2^{(2)}$ integrated dipoles, all of the two-loop Catani poles cancel against the mass factorisation terms and integrated antenna, rendering the entire calculation finite.

Keeping track of precisely what terms were introduced to the subtraction term and precisely where they should be integrated out is much more challenging than at NLO. Fig. 2.1 illustrates the cancellation of terms between the three phase space layers, which has the same general structure regardless of initial state configuration and colour contribution.

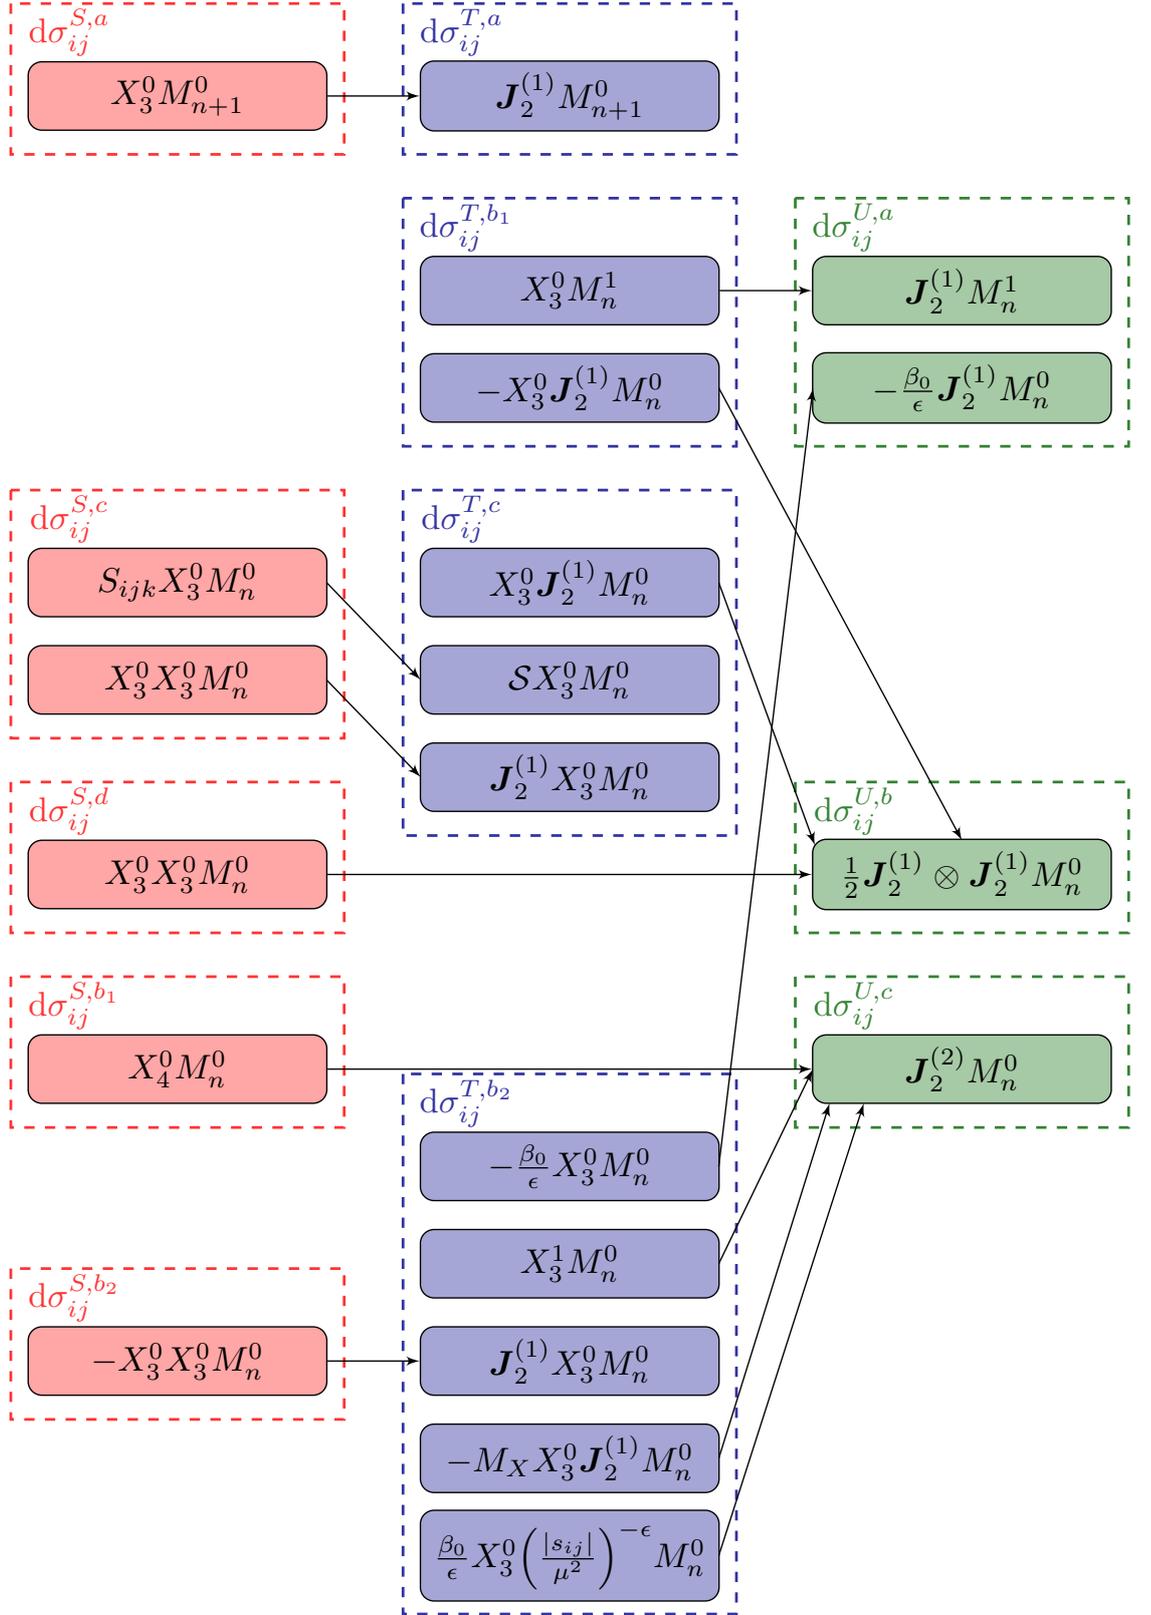


Figure 2.1: A schematic of the contributing terms to the double real (red), real virtual (blue) and double virtual (green) subtraction terms using the antenna formalism. The arrows indicate where terms that are introduced are subsequently integrated out.

2.5 D_4^0 antenna

The implementation of the $D_4^0(i_q, j_g, k_g, l_g)$ antenna in a physical process poses difficulties compared to other antenna because it contains many spurious limits. As discussed previously, the antennae computed from neutralino decays do not precisely have the correct colour decomposition to match the desired QCD radiation pattern. Some of the issues that arise can be dealt with trivially, for example there are two colour orderings within the $D_4^0(i_q, j_g, k_g, l_g)$ antenna [11],

$$(i_q, j_g, k_g, l_g), \quad (2.5.1)$$

$$(i_q, l_g, k_g, j_g), \quad (2.5.2)$$

implying that the $D_4^0(i_q, j_g, k_g, l_g)$ antenna exhibits a line reversal symmetry over the gluon content in the antenna. It is possible to construct unintegrated sub-antenna which have the two colour orderings separated,

$$D_4^0(i_q, j_g, k_g, l_g) = d_4^0(i_q, j_g, k_g, l_g) + d_4^0(i_q, l_g, k_g, j_g), \quad (2.5.3)$$

however the process of repeated partial fractioning to construct the sub-antenna is not symmetric and therefore cannot be integrated analytically [11], implying that only the full $D_4^0(i_q, j_g, k_g, l_g)$ can be used in calculations. When constructing a subtraction term using this antenna, it is therefore required to construct one subtraction term for pairs of colour orderings that are related by a line reversal symmetry over the gluons.

Another issue that arises because of the mismatch in the colour configuration of this antenna is that in a $i_q || j_g || l_g$ triple collinear limit the antenna contains a spurious divergent limit,

$$D_4^0(i_q, j_g, k_g, l_g) \xrightarrow{i_q || j_g || l_g} \tilde{P}_{ijl \rightarrow Q} + \text{angular terms}, \quad (2.5.4)$$

where $\tilde{P}_{ijl \rightarrow Q}$ is a subleading colour splitting function. This oversubtraction can be corrected for by the introduction of new antenna that can regulate the limit. Consider a process with two colour orderings $(i_q, j_g, k_g, l_g, m_{\bar{q}}) + (i_q, l_g, k_g, j_g, m_{\bar{q}})$. Schematically the X_4^0 antenna structure of the subsequent subtraction term would be of the form,

$$+ D_4^0(i_q, j_g, k_g, l_g)$$

$$\begin{aligned}
& + D_4^0(m_{\bar{q}}, j_g, k_g, l_g) \\
& - \tilde{A}_4^0(i_q, j_g, l_g, m_{\bar{q}}),
\end{aligned} \tag{2.5.5}$$

where the $\tilde{A}_4^0(i_q, j_g, l_g, m_{\bar{q}})$ successfully regulates the $\tilde{P}_{ijl \rightarrow Q}$ and $\tilde{P}_{mjl \rightarrow Q}$ divergences from the first and second line respectively. In practice the introduction of the $\tilde{A}_4^0(i_q, j_g, l_g, m_{\bar{q}})$ antenna into a process results in additional spurious single unresolved limits and almost colour connected limits that must again be regulated using the practices described in section 2.4.1.

2.5.1 Gluon initiated D_4^0 antenna

In section 2.1.1, it was shown how to construct the desirable sub-antenna in a final-final configuration, $d_3^0(i_q, j_g, k_g)$, which contains all of the desirable limits through the use of partial fractioning. Constructing an equivalent set of sub-antenna for the D_4^0 antenna in the final-final configuration has been done by exploiting the $\mathcal{N} = 1$ supersymmetric identities of the underlying process [62, 100]. The sub-antenna are formed from many antenna which have each been integrated analytically, hence the requirement that the antenna can be integrated is preserved.

Unfortunately in configurations where a gluon has been crossed into the initial state, constructing sub-antenna in this way is not feasible. This configuration was first relevant for the research presented in Refs. [1, 4] in the dominant qg channel for this process. The overlapping nature of the desired and undesired limits of the antenna makes a construction of sub-antenna difficult. A simpler approach to solving this problem is by introducing new structures to the subtraction terms to regulate the spurious limits.

There are two gluon initiated D_4^0 antennae, $D_4^0(i_q, j_g, \hat{2}_g, k_g)$ and $D_4^0(i_q, \hat{2}_g, j_g, k_g)$. Both of these antennae contain flavour changing and flavour preserving limits, making the reduced matrix element onto which these antenna factorise ambiguous. In both cases the choice is made to factorise onto a flavour preserving matrix element, given that the only other antenna which contains the triple collinear gluon limit $P_{ggg \rightarrow G}$ is the F_4^0 antenna. Introducing an F_4^0 to regulate the flavour preserving limits would require an additional gluon in the process which is not guaranteed.

Consider first the term,

$$+ D_4^0(i_q, j_g, \hat{2}_g, k_g) M_1^0(\hat{1}_q, \hat{2}_g, \widetilde{(ij\bar{k})}_{\bar{q}}) J_1^{(1)}(\{p_1\}), \quad (2.5.6)$$

which is relevant for the qg leading colour contributions to the NNLO corrections to Z +jet production. There are two spurious flavour changing limits in this term, $P_{2ji \leftarrow Q}$ and $P_{\hat{2}ki \leftarrow Q}$, which flavour change the initial state gluon into a quark. This can be regulated by subtracting an A_4^0 containing this spurious limit factorised onto a qg factorised reduced matrix element and adding back the appropriate limit,

$$\begin{aligned} & - A_4^0(\hat{1}, \hat{2}_g, j_g, i_q) M_1^0(\hat{1}_q, \hat{2}_g, \tilde{k}_{\bar{q}}) J_1^{(1)}(\{p_1\}) \\ & + A_4^0(\hat{1}, \hat{2}_g, j_g, i_q) M_1^0(\hat{1}_q, \tilde{k}_g, \hat{2}_{\bar{q}}) J_1^{(1)}(\{p_1\}) \\ & + (j \leftrightarrow k), \end{aligned} \quad (2.5.7)$$

which regulates the spurious triple collinear limits and introduces the correct physical limits. As usual, the single unresolved limits of the newly introduced A_4^0 antennae must also be regulated. These lines are all that are required to regulate the spurious limits of the $D_4^0(i_q, j_g, \hat{2}_g, k_g)$ antenna.

Now consider the configuration,

$$+ D_4^0(i_q, \hat{2}_g, j_g, k_g) M_1^0(\hat{1}_q, \hat{2}_g, \widetilde{(ij\bar{k})}_{\bar{q}}) J_1^{(1)}(\{p_1\}), \quad (2.5.8)$$

where there are similar spurious flavour changing triple collinear limits that can be regulated using the prescription described above. There is, however, a new problematic double collinear limit, $i_q || \hat{2}_g, j_g || k_g$, which is flavour changing. The block of terms arising naturally from $d\sigma_{ij}^{S,b_2}$ regulates this spurious limit,

$$\begin{aligned} & - d_3^0(i, k, j) D_{3,g}^0(\widetilde{(ik)}, \hat{2}, \widetilde{(kj)}) M_1^0(\hat{1}_q, \hat{2}_g, \widetilde{((ik)(kj))}_{\bar{q}}) J_1^{(1)}(\{p_1\}) \\ & - f_{3,g}^0(\hat{2}, j, k) D_{3,g}^0(i, \hat{2}, \widetilde{(jk)}) M_1^0(\hat{1}_q, \hat{2}_g, \widetilde{(ij\bar{k})}_{\bar{q}}) J_1^{(1)}(\{p_1\}), \end{aligned} \quad (2.5.9)$$

however a problem arises in the single collinear $i_q || \hat{2}_g$ limit of the secondary antennae in these terms. In a conventional subtraction term, these limits would cancel against the reduced matrix elements in the $d\sigma_{ij}^{S,a}$ block of terms. The mismatch between the flavour of the factorised partons prevents this cancellation from occurring. The only viable solution is to introduce new antennae that subsequently regulate the single

collinear flavour changing limit of these terms and adds back in the desired physical limits without spoiling other limits in the process. The solution is to introduce the following block of terms,

$$\begin{aligned}
& + f_3^0(\hat{2}, j, k) A_{3,gg \rightarrow qq}^0(\hat{1}, \hat{2}, i) M_1^0(\hat{1}_q, \hat{2}_g, (\widetilde{jk})_{\bar{q}}) J_1^{(1)}(\{p_1\}) \\
& + d_3^0(i, j, k) A_{3,gg \rightarrow qq}^0(\hat{1}, \hat{2}, (\widetilde{ij})) M_1^0(\hat{1}_q, \hat{2}_g, (\widetilde{kj})_{\bar{q}}) J_1^{(1)}(\{p_1\}) \\
& - D_3^0(\hat{1}, j, \hat{2}) A_{3,gg \rightarrow qq}^0(\hat{1}, \hat{2}, i) M_1^0(\hat{1}_q, \hat{2}_g, k_{\bar{q}}) J_1^{(1)}(\{p_1\}) \\
& - A_3^0(\hat{1}, j, i) A_{3,gg \rightarrow qq}^0(\hat{1}_q, \hat{2}_g, (\widetilde{ij})_{\bar{q}}) M_1^0(\hat{1}_q, \hat{2}_g, k_{\bar{q}}) J_1^{(1)}(\{p_1\}) \\
& + 2d_3^0(\hat{1}, j, k) A_{3,gg \rightarrow qq}^0(\hat{1}, \hat{2}, i) M_1^0(\hat{1}_q, \hat{2}_g, (\widetilde{jk})_{\bar{q}}) J_1^{(1)}(\{p_1\}) \\
& + 2A_{3,gg \rightarrow qq}^0(\hat{1}, \hat{2}, i) A_3^0(\hat{1}, j, \hat{2}) M_1^0(\hat{1}_q, \hat{2}_g, k_{\bar{q}}) J_1^{(1)}(\{p_1\}) \\
& - 2A_{3,gg \rightarrow qq}^0(\hat{1}, \hat{2}, i) d_3^0(\hat{2}, j, k) M_1^0(\hat{1}, \hat{2}, (\widetilde{jk})_{\bar{q}}) J_1^{(1)}(\{p_1\}) \\
& - 2A_{3,gg \rightarrow qq}^0(\hat{1}, \hat{2}, i) d_3^0(\hat{1}, j, k) M_1^0(\hat{1}_q, \hat{2}_g, (\widetilde{jk})_{\bar{q}}) J_1^{(1)}(\{p_1\}), \tag{2.5.10}
\end{aligned}$$

which, following symmetrising over the final state gluons j and k , successfully regulates all of the spurious collinear limits introduced. The soft limits must be further corrected by blocks of Eikonal factors which are well understood. The physical limit can be added in simply using,

$$\begin{aligned}
& + f_3^0(\hat{2}, j, k) A_{3,gg \rightarrow qq}^0(\hat{1}, \hat{2}, i) M_1^0(\hat{1}_q, (\widetilde{jk})_g, \hat{2}_{\bar{q}}) J_1^{(1)}(\{p_1\}) \\
& + d_3^0(\hat{1}, j, k) A_{3,gg \rightarrow qq}^0(\hat{1}, \hat{2}, i) M_1^0(\hat{1}_q, (\widetilde{jk})_g, \hat{2}_{\bar{q}}) J_1^{(1)}(\{p_1\}) \\
& - D_3^0(\hat{1}, j, \hat{2}) A_{3,gg \rightarrow qq}^0(\hat{1}, \hat{2}, i) M_1^0(\hat{1}_q, k_g, \hat{2}_{\bar{q}}) J_1^{(1)}(\{p_1\}) \\
& + \text{soft terms.} \tag{2.5.11}
\end{aligned}$$

Clearly the introduction of these additional terms has a significant impact on the construction of the real virtual and double virtual subtraction terms. Nevertheless, it is possible to accommodate the newly introduced terms in such a way that all of the introduced antennae are integrated out and all the singularities are successfully regulated.

2.6 Antenna subtraction with processes involving a vector boson

Matrix elements involving a vector boson have distinctive properties compared to pure QCD processes. These differences have subtle implications for the IR singularity structure of the matrix elements, requiring careful treatment using the antenna formalism.

One significant difference compared to pure QCD processes is that vector bosons must couple to a $q\bar{q}$ pair within a matrix element. IR limits that factorise onto a matrix element that are purely constructed of gluons do not exist. This is a simplification of the calculation compared to pure QCD processes.

Matrix elements involving vector bosons do not exhibit all of the same symmetry properties of other matrix elements (e.g. dijet production, Higgs boson production, etc). As an example, consider the process,

$$e^+(p_5) + e^-(p_6) \rightarrow q(p_1) + g(p_2) + g(p_3) + \bar{q}(p_4), \quad (2.6.1)$$

where for the sake of this discussion it is irrelevant which momenta are in-going or out-going. The tree level amplitude is given by [102],

$$\begin{aligned} \mathcal{A}_4^0(1_q, 2_g, 3_g, 4_{\bar{q}}; 5_{e^+}, 6_{e^-}) &= 2e^2 g_s^2 (-Q^q + v_{L,R}^e v_{L,R}^q \mathcal{P}_z(s_{56})) \\ &\times \sum_{\sigma \in S_2} (T^{a_{\sigma(2)}} T^{a_{\sigma(3)}})_{i_1, i_4} \mathcal{M}_4^0(1_q, \sigma(2)_g, \sigma(3)_g, 4_{\bar{q}}), \end{aligned} \quad (2.6.2)$$

where e and g_s are the EW and QCD coupling factors respectively. Q^q is the charge of the quark flavour q in units of e . $\mathcal{P}_z(s_{56})$ is the ratio of the photon propagator to the Z boson propagator [102],

$$\mathcal{P}_z(s_{56}) = \frac{s_{56}}{s_{56} - M_Z^2 + i\Gamma_Z M_Z}, \quad (2.6.3)$$

where M_Z and Γ_Z are the mass and width of the Z boson respectively. $v_{L,R}^e$ and $v_{L,R}^q$ are the coupling constants for the leptons and quarks respectively. The L and R denote whether the fermions are in the left-handed or right-handed configuration. Explicitly the coupling constants are [102],

$$v_L^e = \frac{-1 + 2 \sin^2 \theta_w}{\sin 2\theta_w}, \quad (2.6.4)$$

$$v_R^e = \frac{2 \sin^2 \theta_w}{\sin 2\theta_w}, \quad (2.6.5)$$

$$v_L^q = \frac{T_3 - 2Q^q \sin^2 \theta_w}{\sin 2\theta_w}, \quad (2.6.6)$$

$$v_R^q = -\frac{2Q^q \sin^2 \theta_w}{\sin 2\theta_w}, \quad (2.6.7)$$

where the T_3 term is +1 for up-type quarks and -1 for down-type quarks.

The important point to note is that the photon couples indiscriminately to the helicity configurations of the fermions, whereas the Z boson couples differently to the fermions depending on their helicities. The Z^μ gauge field contains a superposition of the B^μ and W_3^μ gauge fields, where W_3^μ only couples to the left-handed doublets in the EW Lagrangian.

Now consider an exchange in the momenta of the partons of the form $(p_1, p_2, p_3, p_4) \rightarrow (p_4, p_3, p_2, p_1)$, known as a ‘line reversal’. The colour ordered partial amplitudes, $\mathcal{M}_4^0(1_q, \sigma(2)_g, \sigma(3)_g, 4_{\bar{q}})$, obey a line reversal symmetry [48]. For example,

$$\mathcal{M}_4^0(1_q^+, \sigma(2)_g^+, \sigma(3)_g^-, 4_{\bar{q}}^-) = \mathcal{M}_4^0(4_{\bar{q}}^-, \sigma(3)_g^-, \sigma(2)_g^+, 1_q^+), \quad (2.6.8)$$

hence a pure QCD amplitude is similarly symmetric under line reversal. However, for processes involving a vector boson, the overall coupling on the individual amplitudes differs for the individual helicity configurations, hence the overall amplitude is not symmetric under line reversal of the momentum.

The impact of this on the calculation is that processes involving a quark in the initial state must be treated independently from those with an anti-quark in the initial state. For example, the qg initiated channel is independent from the $g\bar{q}$ calculation. The subtraction terms for the $g\bar{q}$ can be derived from the qg channel by simply performing a line reversal exchange on the reduced matrix elements. For example, a line of the form,

$$- D_3^0(\hat{1}, i, \hat{2}) M_2^0(\hat{1}_q, \hat{2}_g, j_g, k_{\bar{q}}) J_1^{(1)}(\{p_1\}), \quad (2.6.9)$$

in a qg initiated subtraction term would become,

$$- D_3^0(\hat{1}, i, \hat{2}) M_2^0(k_q, j_g, \hat{2}_g, \hat{1}_{\bar{q}}) J_1^{(1)}(\{p_1\}), \quad (2.6.10)$$

in a $g\bar{q}$ subtraction term.

Another implication for processes involving vector bosons is that the matrix elements are dependent on the U(1) charge of the quarks, Q^q . This means that the up-type processes must be treated separately from the down-type processes. This becomes of particular importance for contributions involving two quark pairs with different flavours, for example,

$$q(p_1) + Q(p_2) \rightarrow q(p_3) + Q(p_4) + e^+(p_5) + e^-(p_6), \quad (2.6.11)$$

which is a real correction to the calculation of Z +jet at NLO. This matrix element contains a sum of diagrams over which the Z boson can couple to the $q\bar{q}$ pair or the $Q\bar{Q}$ pair of quarks. As a result of this, there are two initial state collinear singularities, $p_1||p_3$ and $p_2||p_4$, which factorise onto reduced matrix elements of the form $M_3^0(\hat{2}_Q, \hat{1}_g, 4_{\bar{Q}})$ and $M_3^0(\hat{1}_q, \hat{2}_g, 3_{\bar{q}})$ respectively. Knowledge of the flavours must be maintained throughout the calculation.

This becomes particularly challenging when considering a correction of the form,

$$q(p_1) + g(p_2) \rightarrow q(p_3) + Q(p_4) + \bar{Q}(p_5) + e^+(p_6) + e^-(p_7), \quad (2.6.12)$$

which is a double real correction to the calculation of Z +jet at NNLO. This contains many double unresolved singularities which can be classified into three categories based on the flavour of the matrix elements they factorise onto. Either they can factorise onto a matrix element with two quark flavours, for example,

$$M_4^0(\hat{1}_q, \hat{2}_{\bar{Q}}, 3_Q, 4_{\bar{q}}), \quad (2.6.13)$$

or they can factorise onto matrix elements with a single $q\bar{q}$ or $Q\bar{Q}$ pair, for example,

$$M_3^0(\hat{1}_q, \hat{2}_g, 3_{\bar{q}}), \quad (2.6.14)$$

or,

$$M_3^0(\hat{2}_Q, \hat{1}_g, 3_{\bar{Q}}). \quad (2.6.15)$$

In the case of matrix elements involving a single $q\bar{q}$ pair, there is a direct correspondence between the flavour of the quark in the PDF and the flavour of the quarks in the reduced matrix element. However, for limits that factorise into a single $Q\bar{Q}$ pair, the flavour

in the reduced matrix element is decoupled from the flavour of the quark within the PDF. This implies that the corresponding $q\bar{q}$ and $Q\bar{Q}$ terms in the real virtual or double virtual subtraction terms must be treated independently. Their overall factors are different and it would not be appropriate to keep these terms together beyond the double real subtraction terms.

The problems described above make an implementation of the antenna formalism more challenging for processes involving vector bosons. Nevertheless they are not insurmountable. If a subtraction term has been constructed taking into account the issues discussed above, it can *always* be applied to processes that are more symmetric. Antenna subtraction terms are therefore still independent of the colourless particles within the process.

Chapter 3

Drell–Yan phenomenology at the LHC

Drell–Yan production, named after the authors of Ref. [103], is the production of a dilepton pair from a hadron-hadron collision,

$$p + p \rightarrow \gamma^*/Z \rightarrow \ell^+ + \ell^- + X,$$

where the LO production mechanism is depicted in Fig. 3.1. The intermediary propagators are either an off-shell photon or a Z boson and X denotes the hadronic remnants of the event.

The work of Ref. [103] was the first extension of the naïve parton model from deep inelastic scattering (DIS) to a hadron-hadron collider environment. The leading order accurate PDFs could be extracted from DIS measurements and used to make predictions of hadron-hadron collisions [30]. Data from the Alternating Gradient Synchrotron (AGS) gave hints of the possibility of agreement between the naïve parton model and data [104]. Over subsequent years, new experiments were built with higher scattering energies and larger luminosities, revealing the need for more precise theoretical predictions to match the quality of the data. In this chapter precision calculations coming from QCD corrections will be discussed in the context of Drell–Yan production. In particular, the importance of precision predictions of Drell–Yan production in association with hadronic jets,

$$p + p \rightarrow \gamma^*/Z \rightarrow \ell^+ + \ell^- + \text{Jet}(s) + X,$$

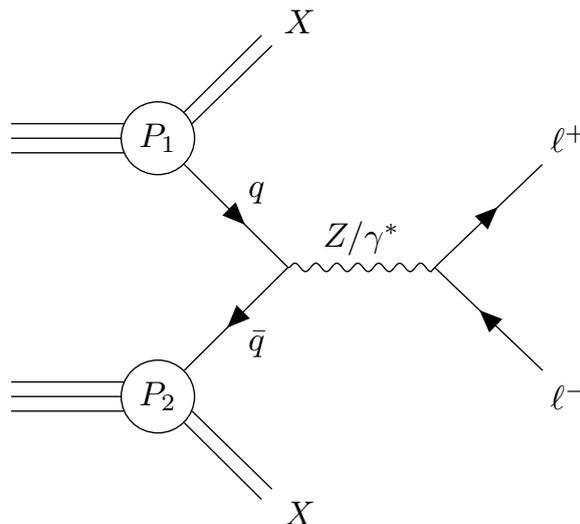


Figure 3.1: The leading order contribution to the Drell–Yan cross section in the naïve parton model. X denotes the hadronic remnants of the event.

will be motivated.

Henceforth Drell–Yan production in association with hadronic jets will simply be referred to as ‘ Z +jet(s)’. Nevertheless, the contribution coming from the photon propagator is included implicitly. Furthermore, unless explicitly stated otherwise, we are always inclusive in the number of final state jets (e.g. Z +jet production implies Z boson production with *at least* one jet).

3.1 Motivation for precision predictions of Drell–Yan observables

The motivation for precision Drell–Yan predictions is driven by the large volume of data available at both the LHC and previous colliders. The Z resonance produces a large event rate at the LHC and is the dominant contribution to the dilepton cross section. As Tab. 3.1 demonstrates however, the primary decay modes of the Z boson are hadronic and invisible final states. Despite this, leptonic signals are the easiest to reconstruct and hence the best decay mode to perform precision phenomenology¹. It

¹Often when performing measurements the $Z \rightarrow \tau^+\tau^-$ decay is not considered part of the Drell–Yan signature production mechanism. This is because τ lepton has an extremely short lifetime of $(290.3 \pm 0.5 \times 10^{-15})$ s [25] and its mass is sufficiently large that it decays into either hadronic or leptonic final states. Great care must be taken to reconstruct τ particles and hence for the purposes of

| Decay Type | Branching Ratio |
|----------------------------------|-----------------------|
| $Z \rightarrow e^+e^-$ | $(3.363 \pm 0.004)\%$ |
| $Z \rightarrow \mu^+\mu^-$ | $(3.366 \pm 0.007)\%$ |
| $Z \rightarrow \tau^+\tau^-$ | $(3.370 \pm 0.008)\%$ |
| $Z \rightarrow \text{invisible}$ | $(20.0 \pm 0.06)\%$ |
| $Z \rightarrow \text{hadrons}$ | $(69.91 \pm 0.06)\%$ |

Table 3.1: The branching ratios for the dominant decay modes of the Z boson [25].

is for this reason that Drell–Yan events are used to constrain quark PDFs over a wide range of x and Q^2 [105].

The gluon PDF can be constrained by including data involving hadronic final states, requiring knowledge of theoretical corrections to processes involving well resolved final state jets. Typically PDFs are described up to a specific order in a perturbative expansion of α_s . To use a specific set of data in a PDF fit, the corresponding partonic cross section must be known up to the desired order for the fitting. This is a strong motivation for calculating the inclusive Z +jet cross section at NNLO. This calculation can then be used to create NNLO accurate PDF sets to constrain the gluon PDF over a wide range of x and Q^2 .

Drell–Yan processes also pose a significant background to Beyond the Standard Model (BSM) searches. Typical searches for supersymmetry, for example, involve searches for isolated opposite sign same flavour leptons in association with hard jets [106, 107]. Drell–Yan processes and associated jets are therefore an important background for this signal of BSM physics. Clearly the SM backgrounds must be described as precisely as possible to resolve supersymmetric signals which are going to be extremely small over a large SM background. This is of particular importance at high scattering energies for particular observables where production rates for TeV scale BSM physics might be detected over SM backgrounds.

In principle, a precision calculation of Drell–Yan production can also be used for providing predictions for the invisible decay modes of the Z boson. This is of particular importance for dark matter searches, where if a dark matter particle is produced at the

precision phenomenology the electron and muon decay modes are far more constraining measurements.

LHC it will not be directly detectable. The SM backgrounds to this are processes that produce neutrinos, which similarly pass through the detector and appear as missing transverse momentum.

3.2 Theoretical corrections to Drell–Yan production

Since the work of Ref. [103], a variety of precision corrections to the Drell–Yan process have been considered. The NLO QCD corrections to the Drell–Yan cross section were calculated a long time ago by several groups [108–114] and is considered to be a benchmark calculation of SM phenomenology. The calculations of Refs. [108–114] marked the first use of the improved parton model (described in section 1.7.1) within a hadron-hadron collider environment.

The NLO K factor, defined as,

$$K = \frac{\sigma_{\text{NLO}}}{\sigma_{\text{LO}}}, \quad (3.2.1)$$

was found to be very large; typically in the region of $K \sim 0.5$ for LHC scattering energies using representative experimental cuts. Despite the potential for the perturbative cross section to be poorly converging, there is an intuitive argument for why the NLO corrections are significant. As Fig. 3.2(a) shows, the NLO corrections introduce a new gg initiated channel. As discussed in section 1.7, at small momentum fractions the gluon PDF is much larger than all of the quark PDFs for typical LHC scattering energies. If there is sufficient energy to readily produce Z bosons from a combination of partons including an initial state gluon then these channels will receive a large weighting factor from the gluon PDF compared to purely quark initiated processes.

Fig. 3.2(b) demonstrates the possibility of a new gg initiated partonic channel appearing first at NNLO. Here two factors prevent the gg channel from being large,

- to produce the Z boson from two initial state gluons requires that they both must have a relatively large momentum fraction at typical scattering energies for the LHC, which is where the gluon PDFs are small.
- It is suppressed by an additional power of α_s .

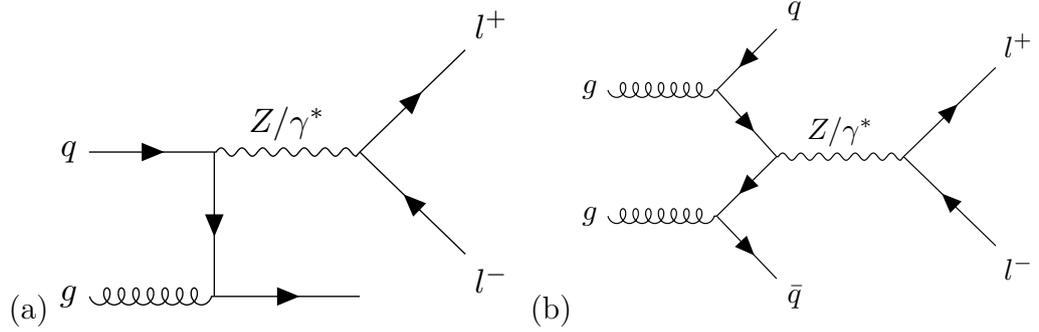


Figure 3.2: Feynman diagrams depicting (a) a real correction to Drell–Yan production from a $q\bar{q}$ initiated channel at $\mathcal{O}(\alpha_s)$ and (b) a double real correction to Drell–Yan production from a $g\bar{g}$ initiated channel at $\mathcal{O}(\alpha_s^2)$.

Subsequently the NNLO corrections to the Drell–Yan process were calculated [115, 116] and indeed the NNLO K factor, defined as,

$$K = \frac{\sigma_{\text{NNLO}}}{\sigma_{\text{NLO}}}, \quad (3.2.2)$$

was found to be significantly smaller than the NLO K factor, roughly $K \sim 0.01$ at LHC scattering energies. This suggests there is a good convergence of the perturbative expansion and that the $g\bar{g}$ initiated channel is overall a small contribution to the total cross section. It is a common feature of many processes that the NLO K factor is large as new partonic channels are revealed and the NNLO corrections stabilise the calculation. Nevertheless, the NNLO calculation was vital in reducing the scale uncertainty with respect to the NLO calculation, allowing for precision phenomenology to be performed.

Many other forms of corrections beyond fixed order QCD have also been presented for Drell–Yan production. Indeed, the NLO EW corrections $\mathcal{O}(\alpha)$ have been known for some time [117]. More recently, the EW corrections have been matched to a parton shower (PS) [118], allowing for theoretical predictions in phase space configurations where the fixed order perturbative expansion is ill-defined. The NNLO fixed order corrections have also been combined with a resummation of next-to-next-to-leading logarithmic effects [119, 120] to describe the Z boson at small transverse momentum. Very recently mixed QCD-EW corrections $\mathcal{O}(\alpha_s\alpha)$ were also considered [121, 122].

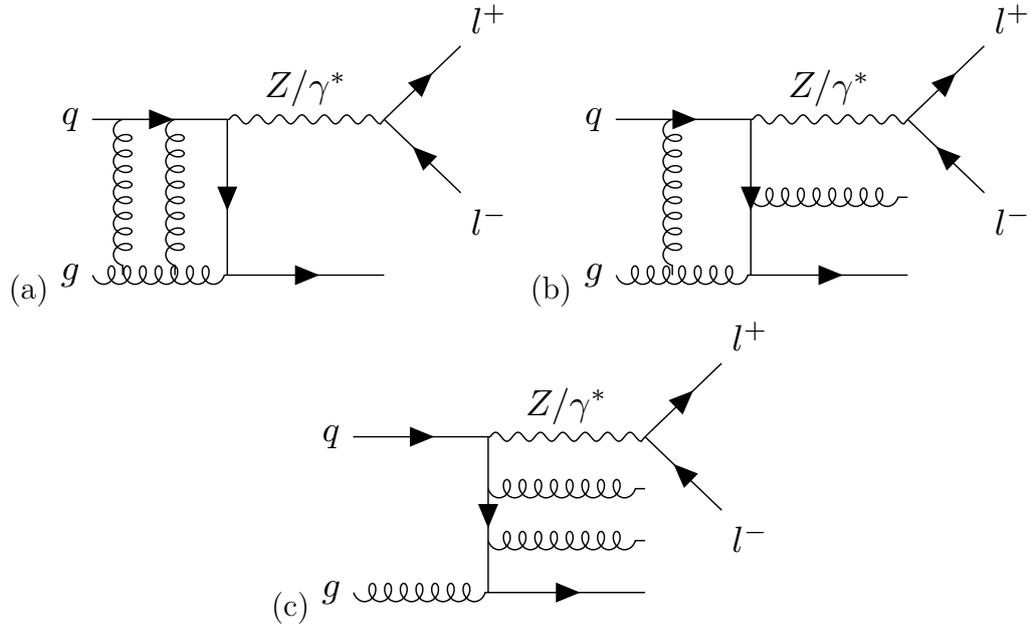


Figure 3.3: Representative Feynman diagrams for (a) two-loop Z boson plus three parton amplitudes (b) one-loop Z boson plus four parton amplitudes and (c) tree level Z boson plus five parton amplitudes.

3.3 Drell–Yan production in association with hadronic jets

Drell–Yan production in association with hadronic jets has also been intensively studied. The NLO QCD corrections for Z +jet [123], Z +2 jets [124, 125], Z +3 jets [126] and Z +4 jets [127] are known while the NLO EW corrections for Z +jet [128, 129] and Z +2 jets [130] have also been derived.

The work in this thesis relates to a new calculation of the NNLO contributions to the neutral current Drell–Yan process in which the dilepton pair is produced in association with a hard, visible hadronic jet. Representative Feynman diagrams contributing to this calculation are shown in Fig. 3.3. The results are obtained from a parton level event generator NNLOJET outlined in chapter 4 in a fully differential form, including the Z/γ^* boson decay to two charged leptons. The final state of the hard scattering process is completely reconstructable and the application of an invariant mass cut on the lepton pair can ensure that the process is dominated by resonant Z bosons.

The NNLO corrections to Z +jet production in hadronic collisions receive contributions from three types of parton-level processes: (a) the two-loop corrections to Z

boson-plus-three-parton processes [131–134], (b) the one-loop corrections to Z boson-plus-four-parton processes [102, 124, 125, 135–137] and (c) the tree level Z boson-plus-five-parton processes [124, 125, 138–140]. The subtraction terms employed to regulate the IR singularity structure were regulated using the antenna formalism outlined in chapter 2.

Following the calculation presented in Ref. [1], a second calculation of Z +jet production at NNLO precision has been presented and published [141]. In coordination with the authors of [141], an in-depth comparison was performed. Using the same setup (cuts, parton distributions, scale choice). This comparison uncovered an error in the numerical code used in [141], which alters their published results. After correction of this error, the code developed in [141] agrees with the calculation outlined in this section.

3.4 Transverse momentum distribution of the Z boson

As illustrated in Fig. 3.4, the transverse momentum of the Z boson, p_T^Z , is generated by the emission of QCD radiation so that the fixed order calculation at $\mathcal{O}(\alpha_s^2)$, which is NNLO for the inclusive cross section, corresponds to only an NLO accurate prediction of the transverse momentum distribution.

The calculation presented in the previous section can be used to compute the transverse momentum distribution at NNLO precision. To achieve this, the requirement of observing a final state jet is relaxed. Instead a low transverse momentum cut on the Z boson is imposed to ensure the IR finiteness of the NNLO calculation. This enforces the presence of final state partons to balance the transverse momentum of the Z boson. The production of Z bosons (or, more generally, of lepton pairs with given invariant mass) at large transverse momentum has been studied extensively at the LHC by the ATLAS [142, 143], CMS [144, 145] and LHCb [146] experiments. ATLAS and CMS both observed a tension between their measurements and existing NLO QCD predictions, highlighting the potential importance of higher order corrections to this process.

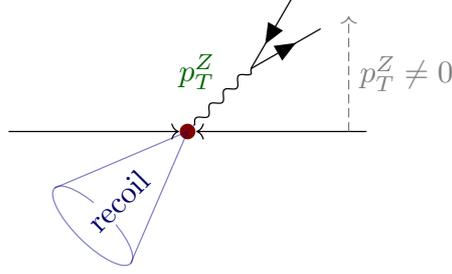


Figure 3.4: A schematic diagram demonstrating the Z boson recoiling against hard radiation.

In order to reduce the systematic uncertainty on the measurement, the transverse momentum distribution is commonly normalised to the p_T inclusive Z boson production cross section. To make a comparison to both the unnormalised and normalised data sets from ATLAS and CMS, both the absolute (unnormalised) differential cross section,

$$\left. \frac{d\sigma}{dp_T^Z} \right|_{p_T^Z > 20 \text{ GeV}} \equiv \frac{d\sigma_{LO}^{ZJ}}{dp_T^Z} + \frac{d\sigma_{NLO}^{ZJ}}{dp_T^Z} + \frac{d\sigma_{NNLO}^{ZJ}}{dp_T^Z} \quad (3.4.1)$$

and the differential cross section normalised to the relevant Drell–Yan cross section,

$$\frac{1}{\sigma} \cdot \left. \frac{d\sigma}{dp_T^Z} \right|_{p_T^Z > 20 \text{ GeV}}$$

with

$$\sigma = \int_0^\infty \frac{d\sigma}{dp_T^Z} dp_T^Z \equiv \sigma_{LO}^Z + \sigma_{NLO}^Z + \sigma_{NNLO}^Z, \quad (3.4.2)$$

are calculated.

Both experiments present their measurements in the form of fiducial cross sections for a restricted kinematical range of the final state leptons (in invariant mass, transverse momentum and rapidity). In view of a comparison between data and theory, this form of presenting the experimental data is preferable over a cross section that is fully inclusive in the lepton kinematics (requiring a theory based extrapolation into phase space regions outside the detector coverage). Consequently, the theoretical calculation must take proper account of these restrictions in the final state lepton kinematics.

The unnormalised p_T^Z distribution represents an absolute cross section measurement based on event counting rates. As with any absolute measurement, it has the disadvantage of being sensitive to the proper modelling of acceptance corrections, and of relying on the absolute determination of the integrated luminosity of the data sample

under consideration. At the LHC the luminosity uncertainty alone amounts to about 3%. In order to reduce the luminosity uncertainty, the data can be normalised to the Drell–Yan cross section for the corresponding fiducial phase space. This is obtained from the cross section for Z boson production with the same transverse momentum and rapidity cuts on the individual leptons, but integrated over all possible transverse momenta of the Z boson.

Chapter 4

Numerical implementation of NNLO calculations

To make predictions of interest for particle physics phenomenology, numerical integration techniques are mandatory. Cross section calculations and differential distributions require an integral over a phase space, $d\sigma_n$, which have no known analytical solution for non-trivial processes. For a typical collider final state the phase space can be of extremely high dimensionality.

For numerical integrals over a large number of dimensions, Monte Carlo sampling is the most efficient technique for evaluating the integral. The focus of this chapter will be on the construction of the parton level event generator NNLOJET, which uses the antenna subtraction method to calculate observables up to NNLO as a perturbative expansion in α_s . Sections 4.1, 4.3, 4.4 and 4.5 will discuss the methods used to validate the predictions made in the subsequent chapters. This should not be considered as an exhaustive list of the checks performed on the numerical calculation however. Further details of the numerical checks can be found in appendix B. Sections 4.2, 4.6 and 4.7 will instead focus on the construction of the numerical code and the significant optimisations required to make the code as efficient as possible.

4.1 Numerical testing of subtraction terms

As Fig. 2.1 illustrates, a subtraction term contains many components. Extensive testing is therefore required before it can be integrated numerically. In the case of the real

virtual and double virtual subtraction terms, the explicit cancellation of the ϵ pole structure can be checked analytically. The subtraction term can be imported into a larger program written in FORM [147] and, using purely algebraic techniques, the ϵ pole structure can be checked against the known singularity structure in terms of Catani operators.

In the case of implicit singularities from phase space integrals it is simpler to test the limits numerically. Given that a functioning subtraction term must regulate *all* of the divergences of the matrix element, every limit of the subtraction term has to be probed to ensure that the divergences cancel. To do this the phase space generator RAMBO [148] is used, which generates phase space points with a flat weighting. The fact that events are generated in an unbiased way is an important requirement of the testing procedure. All of the unresolved configurations relevant for a particular integral must be tested, even if during Monte Carlo integration the likelihood of sampling such a configuration is small.

RAMBO can generate phase spaces for massless and massive particles up to extremely high multiplicities. Whilst it was not specifically designed to generate unresolved phase space configurations, it is possible to force it to generate these phase spaces through careful manipulation of the masses of intermediate decaying particles. For example, consider the phase space generated by,

$$p_1 + p_2 \rightarrow p_3 + p_{45} \rightarrow p_3 + p_4 + p_5, \quad (4.1.1)$$

where p_{45} is off-shell and has a mass of $\sqrt{s_{45}}$ and all other momenta are on-shell. The scattering energy of the process is denoted by \sqrt{s} .

In the limit $\sqrt{s_{45}} \rightarrow \sqrt{s}$, the momenta of p_3 is driven soft. This can be controlled numerically by setting,

$$s_{45} = s \cdot (1 - x_i), \quad (4.1.2)$$

where x_i is small dimensionless parameter. x_i controls the distance the phase space point is away from the absolute p_3 soft limit.

Using a similar approach, in the limit $\sqrt{s_{45}} \rightarrow 0$, the momenta p_4 and p_5 become collinear. Again this can be controlled by an x_i parameter,

$$s_{45} = s \cdot x_i. \quad (4.1.3)$$

This method of controlling the proximity to the unresolved limit can also be implemented for initial-final collinear limits. This is performed by generating a final-final collinear limit and rotating one of the final state partons into the initial state and boosting the configuration such that both initial state partons are in their centre-of-mass frame.

Using these ingredients, the generation of double unresolved configurations from an extended phase space with an additional parton are also straightforward. In the case of double collinear limits, the phase space can be generated from,

$$p_1 + p_2 \rightarrow p_{34} + p_{56} \rightarrow p_3 + p_4 + p_5 + p_6, \quad (4.1.4)$$

where p_{34} and p_{56} have masses $\sqrt{s_{34}}$ and $\sqrt{s_{56}}$ respectively. In a double collinear limit $\sqrt{s_{34}} \rightarrow 0$ and $\sqrt{s_{56}} \rightarrow 0$ simultaneously. All of the other double unresolved limits can be generated using an iterated approach,

$$p_1 + p_2 \rightarrow p_3 + p_{456} \rightarrow p_3 + p_4 + p_{56} \rightarrow p_3 + p_4 + p_5 + p_6, \quad (4.1.5)$$

where p_{456} has a mass of $\sqrt{s_{456}}$ and p_{56} has a mass of $\sqrt{s_{56}}$. The soft collinear limit can be generated by taking $\sqrt{s_{456}} \rightarrow \sqrt{s}$ (making p_3 soft) and $\sqrt{s_{56}} \rightarrow 0$ such that p_5 is collinear with p_6 . Similarly to the single collinear limit, for a triple collinear limit $\sqrt{s_{456}} \rightarrow 0$.

Given this simple implementation for formulating single and double unresolved phase space configurations, it is now straightforward to use these to test whether the subtraction term is actually working correctly. A ratio R can be defined such that,

$$R = \frac{d\sigma_{ij}^R}{d\sigma_{ij}^{S,\text{NLO}}}, \quad (4.1.6)$$

where ideally in all of the appropriate single unresolved limits $R \rightarrow 1$. Similar ratios can be defined for the real virtual and double real contributions.

In collinear limits where a gluon splits either into a pair of gluons or a quark anti-quark pair, the ratio is spoiled by the presence of the azimuthal correlation terms discussed in section 1.6. A simple solution, proposed in Refs. [100, 101], is to perform a rotation of $\pi/2$ about the shared collinear axis and evaluating the matrix element and subtraction term again. The combination of these two phase space evaluations

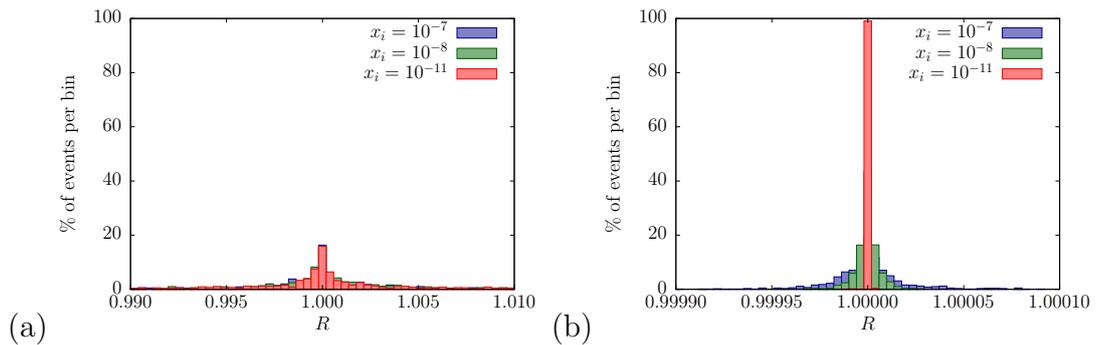


Figure 4.1: Example spike plots for the triple gluon collinear limit leading colour contribution to the $q\bar{q} \rightarrow ggg$ for Z +jet at NNLO. x_i is a small dimensionless parameter that controls the proximity of the phase space to the triple collinear limit. (a) shows the limit without an azimuthal rotation and (b) shows the limit with an azimuthal rotation. The x axis differs between the two plots for clarity.

cancels the angular dependent terms, and hence the $R \rightarrow 1$ condition required for a functional subtraction term is restored.

This approach can also be extended to double unresolved limits. In a double collinear limit four phase space evaluations, two phase space evaluations about each collinear axis are sufficient to cancel the azimuthal correlations. In triple collinear limits a $\pi/2$ rotation around the shared triple collinear axis will cancel all of the azimuthal terms. The use of these rotations to cancel azimuthal terms implies that the antenna subtraction scheme is not a fully local subtraction method. Antenna subtraction falls under the category of being ‘quasi-local’, since the leading singularities are cancelled locally and there is an established prescription for constructing rotated phase spaces that can cancel the angular terms.

Fig. 4.1(a) is a spike plot showing the ratio R in a triple collinear gluon limit for the leading colour channel $q\bar{q} \rightarrow ggg$ for Z +jet at NNLO with the azimuthal rotation turned off. Fig. 4.1(b) demonstrates the same limit with an added azimuthal rotation. It is clear that once the azimuthal rotation is applied, the distribution of R values is significantly closer to 1 for a given value of x_i .

An important point to note is that for double real subtraction terms it is not sufficient to test only the double unresolved configurations. Single unresolved configurations must be tested independently, even though they are subleading singularities with respect to double unresolved limits. The internal cancellation within the subtraction term for single unresolved configurations is very different from double unresolved

limits, meaning it is plausible that all of the double unresolved limits work correctly but the single unresolved limits do not function correctly.

4.2 Virtual numerical integration

The implementation of processes involving at least one-loop requires some special treatment compared to processes involving purely tree level matrix elements. One-loop matrix elements are constructed as d -dimensional objects to allow the UV and IR singularities to be written as a series of poles in ϵ . This approach is well grounded as an analytical expression defined in d -dimensions, however any numerical evaluations of the matrix element must be performed in 4-dimensions. Clearly for the numerical integral to be well defined, the finite remainder of the integrand must be well understood as $\epsilon \rightarrow 0$.

As shown in Eq. (1.6.1), a one-loop matrix element can be written as a series of Catani operators acting on the tree level matrix element, plus a finite remainder contribution. There is, however, an ambiguity in Eq. (1.6.1) as to the dimensionality of the overall expression. The tree level matrix elements, like their one-loop counterparts, can be defined in d -dimensions as a finite series in ϵ ,

$$|\mathcal{M}_n^{(0)}(\mu_R^2, \{p_i\})\rangle_{\dim=d} = \sum_{i=0} \epsilon^i |\mathcal{M}_{n,i}^{(0)}(\mu_R^2, \{p_i\})\rangle, \quad (4.2.1)$$

where in the limit $\epsilon = 0$, the 4-dimensional result must be recovered,

$$|\mathcal{M}_n^{(0)}(\mu_R^2, \{p_i\})\rangle_{\dim=4} = |\mathcal{M}_{n,0}^{(0)}(\mu_R^2, \{p_i\})\rangle. \quad (4.2.2)$$

If the expression in Eq. (1.6.1) is taken to be in d -dimensions, then the coefficients of the ϵ expansion of Eq. (4.2.1) generate both new pole structures and additional finite remainder terms. These ultimately have no impact on the final physical result however because the choice of the dimensionality of the tree level contribution must match between the one-loop matrix element and the subtraction term. Any additional structures from the higher order ϵ contributions must ultimately cancel. This means that there is a choice as to when to set the limit $\epsilon = 0$ when performing virtual integration,

- perform the cancellation between the matrix element and subtraction term analytically in d -dimensions and then set the limit $\epsilon = 0$ in the resulting expression to be evaluated numerically,
- take all of the tree level matrix elements to be in 4-dimensions from the beginning and implicitly assume that higher order d -dimensional terms from the tree level matrix elements must cancel between the one-loop matrix element and the subtraction term. This is effectively the HV scheme choice.

Both methods should arrive back at the same numerical result. Typically though it is simpler to implement the second option and therefore this is the adopted approach in the NNLOJET program for both virtual and real virtual integrals.

In the case of double virtual integration, a similar choice of the dimensionality must be made. As Eq. (1.6.18) shows however, there is an added complication that there is a pole structure that factorises onto a reduced one-loop contribution, $\mathbf{I}_n^{(1)}(\epsilon; \mu_R^2, \{p_i\}) |\mathcal{M}_n^{(1)}(\mu_R^2, \{p_i\})\rangle$. Here it becomes ambiguous to define a 4-dimensional limit of the one-loop matrix element since it is a divergent quantity in the $\epsilon \rightarrow 0$ limit. It is simpler in the case of the two-loop matrix elements to keep them entirely in d -dimensions.

In principle, this would mean that the double virtual integration must be performed using option one above. This would require an algebraic expression for the two-loop matrix element which is understood well enough that it can be combined with the subtraction term before a numerical implementation. This is not ideal since there is a significant degree of complexity in constructing two-loop matrix elements for non-trivial processes. To overcome this challenge, a simple approach is to define a Catani subtracted two-loop matrix element in d -dimensions,

$$d\sigma_{ij}^{VV,CS} = \left(d\sigma_{ij}^{VV} - d\sigma_{ij}^C \right) \Big|_{\dim=d}, \quad (4.2.3)$$

where $d\sigma_{ij}^C$ denotes the Catani pole structure Eq. (1.6.22) integrated over an appropriate phase space. $d\sigma_{ij}^{VV,CS}$ is clearly a finite quantity in the limit $\epsilon \rightarrow 0$ and is completely well defined. Similarly one can also formulate a Catani subtracted definition of the subtraction term,

$$d\sigma_{ij}^{U,CS} = \left(d\sigma_{ij}^U - d\sigma_{ij}^C \right) \Big|_{\dim=4}, \quad (4.2.4)$$

where $d\sigma_{ij}^C|_{\dim=4}$ is defined such that all of the tree level matrix elements are in 4-dimensions, including the tree level contributions to the one-loop matrix elements within the Catani pole structure. This allows the double virtual integral to be formulated as,

$$\int_{d\sigma_n} [d\hat{\sigma}_{ij}^{VV} - d\hat{\sigma}_{ij}^U] = \int_{d\sigma_n} (d\sigma_{ij}^{VV} - d\sigma_{ij}^C)|_{\dim=d} - \int_{d\sigma_n} (d\sigma_{ij}^U - d\sigma_{ij}^C)|_{\dim=4}, \quad (4.2.5)$$

or,

$$\int_{d\sigma_n} [d\hat{\sigma}_{ij}^{VV} - d\hat{\sigma}_{ij}^U] \stackrel{\epsilon \rightarrow 0}{=} \int_{d\sigma_n} [d\hat{\sigma}_{ij}^{VV,CS} - d\hat{\sigma}_{ij}^{U,CS}]. \quad (4.2.6)$$

Using this formalism, the two-loop matrix element can be imported into the NNLOJET code with relative ease. There are no prerequisites on the structure or the method used to compute the two-loop matrix element, only that they are Catani subtracted. The two-loop matrix element does not even necessarily need to be colour decomposed, although this is a desirable feature to maintain.

Once the Catani subtracted definition of the subtraction term has been defined for a given parton multiplicity and flavour configuration, simply changing the tree level and one-loop matrix elements to a new process with the same parton configuration becomes a trivial exercise. Using this approach, the two-loop integration is as flexible as possible both for the construction and testing of the subtraction terms and the two-loop matrix elements.

4.2.1 Plus distributions

Virtual integration routines require the use of integrated antenna functions, as outlined in section 2.3.2. The most generalised integrated antenna function is given as a function of the form,

$$\begin{aligned} \mathcal{X}(x_1, x_2) &= A(x_1, x_2)\delta(1-x_1)\delta(1-x_2) \\ &+ B(x_1, x_2)\delta(1-x_1) \\ &+ \sum_i^n \mathcal{D}_i(x_2)C_i(x_1, x_2)\delta(1-x_1) \\ &+ D(x_1, x_2)\delta(1-x_2) \end{aligned}$$

$$\begin{aligned}
& + \sum_i^n \mathcal{D}_i(x_1) E_i(x_1, x_2) \delta(1 - x_2) \\
& + F(x_1, x_2) \\
& + \sum_i^n \mathcal{D}_i(x_1) G_i(x_1, x_2) \\
& + \sum_i^n \mathcal{D}_i(x_2) H_i(x_1, x_2) \\
& + \sum_i^n \sum_j^n \mathcal{D}_i(x_1) \mathcal{D}_j(x_2) I_{ij}(x_1, x_2), \tag{4.2.7}
\end{aligned}$$

where the set of functions $A(x_1, x_2), \dots, I(x_1, x_2)$ are smooth polynomials in terms of x_1 and x_2 . The functions $\mathcal{D}_i(x_1)$ are a series of plus distributions,

$$\mathcal{D}_i(x) = \left(\frac{\log^i(1-x)}{1-x} \right)_+ . \tag{4.2.8}$$

The maximum possible value for n is one for tree level three parton integrated antenna functions. Higher order plus distributions only appear in the NNLO antenna functions.

There are two equivalent definitions of the plus distribution. The first definition is given by considering its integral with respect to a smooth test function $g(x)$ [30],

$$\int_0^1 dx g(x) f(x)_+ = \int_0^1 dx [g(x) - g(1)] f(x), \tag{4.2.9}$$

which regulates the singularity of $f(x)$ as $x \rightarrow 1$. The second definition is given by,

$$f(x)_+ = f(x) - \delta(1-x) \int_0^1 dx' f(x'), \tag{4.2.10}$$

where a substitution of this relation into Eq. (4.2.9) reveals that they are equivalent definitions.

From Eq. (4.2.7), there are four integration regions for the specific combinations of δ -functions,

1. $\delta(1-x_1)\delta(1-x_2)$,
2. $\delta(1-x_1)$,
3. $\delta(1-x_2)$,
4. no δ -functions.

There is a choice in terms of how to integrate these regions numerically. The first method is to rescale the phase space for the four regions,

$$\begin{aligned}
d\sigma^T = & -\mathcal{N}_{\text{NLO}}^V \mathcal{C}(N, n_f) \sum_{ij} \int_0^1 \frac{dz_1}{z_1} \int_0^1 \frac{dz_2}{z_2} \frac{1}{S_n} \left\{ \right. \\
& + \int_0^1 dx_1 \int_0^1 dx_2 \left[f_i(z_1, \mu_F^2(\{p_n\})) f_j(z_2, \mu_F^2(\{p_n\})) \right. \\
& \quad \times d\Phi_n(p_3, \dots, p_{n+2}; z_1 p_1, z_2 p_2) \mathcal{X}_1(x_1, x_2) \\
& \quad \left. \times M_{n+2}^0(\{z_1 z_2 p_{n+2}\}) J_n^{(n)}(\{p_n\}) \right] \\
& + \int_0^1 dx_1 \int_0^1 \frac{dx_2}{x_2} \left[f_i(z_1, \mu_F^2(\{x_2 p_n\})) f_j(z_2, \mu_F^2(\{x_2 p_n\})) \right. \\
& \quad \times d\Phi_n(p_3, \dots, p_{n+2}; z_1 p_1, x_2 z_2 p_2) \mathcal{X}_2(x_1, x_2) \\
& \quad \left. \times M_{n+2}^0(\{x_2 z_1 z_2 p_{n+2}\}) J_n^{(n)}(\{x_2 p_n\}) \right] \\
& + \int_0^1 \frac{dx_1}{x_1} \int_0^1 dx_2 \left[f_i(z_1, \mu_F^2(\{x_1 p_n\})) f_j(z_2, \mu_F^2(\{x_1 p_n\})) \right. \\
& \quad \times d\Phi_n(p_3, \dots, p_{n+2}; x_1 z_1 p_1, z_2 p_2) \mathcal{X}_3(x_1, x_2) \\
& \quad \left. \times M_{n+2}^0(\{x_1 z_1 z_2 p_{n+2}\}) J_n^{(n)}(\{x_1 p_n\}) \right] \\
& + \int_0^1 \frac{dx_1}{x_1} \int_0^1 \frac{dx_2}{x_2} \left[f_i(z_1, \mu_F^2(\{x_1 x_2 p_n\})) f_j(z_2, \mu_F^2(\{x_1 x_2 p_n\})) \right. \\
& \quad \times d\Phi_n(p_3, \dots, p_{n+2}; x_1 z_1 p_1, x_2 z_2 p_2) \mathcal{X}_4(x_1, x_2) \\
& \quad \left. \times M_{n+2}^0(\{x_1 x_2 z_1 z_2 p_{n+2}\}) J_n^{(n)}(\{x_1 x_2 p_n\}) \right] \left. \right\}, \tag{4.2.11}
\end{aligned}$$

where the functions $\mathcal{X}_1(x_1, x_2), \dots, \mathcal{X}_4(x_1, x_2)$ are functions that constitute the full antenna separated into the four regions,

$$\begin{aligned}
\int_0^1 \frac{dx_1}{x_1} \int_0^1 \frac{dx_2}{x_2} \mathcal{X}(x_1, x_2) = & \int_0^1 \frac{dx_1}{x_1} \int_0^1 \frac{dx_2}{x_2} \left[\mathcal{X}_1(x_1, x_2) + \mathcal{X}_2(x_1, x_2) \right. \\
& \left. + \mathcal{X}_3(x_1, x_2) + \mathcal{X}_4(x_1, x_2) \right], \tag{4.2.12}
\end{aligned}$$

where,

$$\begin{aligned}
\mathcal{X}_1(x_1, x_2) = & A(1, 1) - \sum_i^n \frac{\log^i(1-x_2)}{1-x_2} C_i(1, 1) - \sum_i^n \frac{\log^i(1-x_2)}{1-x_2} E_i(1, 1) \\
& + \sum_i^n \sum_j^n \frac{\log^i(1-x_1) \log^j(1-x_2)}{1-x_1} I_{ij}(1, 1), \tag{4.2.13}
\end{aligned}$$

$$\mathcal{X}_2(x_1, x_2) = B(1, x_2) + \sum_i^n \frac{\log^i(1-x_2)}{1-x_2} C_i(1, x_2) - \sum_i^n \frac{\log^i(1-x_1)}{1-x_1} G_i(1, x_2)$$

$$- \sum_i^n \sum_j^n \frac{\log^i(1-x_1)}{1-x_1} \frac{\log^j(1-x_2)}{1-x_2} I_{ij}(1, x_2), \quad (4.2.14)$$

$$\begin{aligned} \mathcal{X}_3(x_1, x_2) = & D(x_1, 1) + \sum_i^n \frac{\log^i(1-x_1)}{1-x_1} E_i(x_1, 1) - \sum_i^n \frac{\log^i(1-x_1)}{1-x_1} H_i(x_1, 1) \\ & - \sum_i^n \sum_j^n \frac{\log^i(1-x_1)}{1-x_1} \frac{\log^j(1-x_2)}{1-x_2} I_{ij}(x_1, 1), \end{aligned} \quad (4.2.15)$$

$$\begin{aligned} \mathcal{X}_4(x_1, x_2) = & F(x_1, x_2) + \sum_i^n \frac{\log^i(1-x_1)}{1-x_1} G_i(x_1, x_2) + \sum_i^n \frac{\log^i(1-x_1)}{1-x_1} H_i(x_1, x_2) \\ & + \sum_i^n \sum_j^n \frac{\log^i(1-x_1)}{1-x_1} \frac{\log^j(1-x_2)}{1-x_2} I_{ij}(x_1, x_2). \end{aligned} \quad (4.2.16)$$

This method of integrating the regions is not ideal. As Eq. (4.2.11) shows, for a given set of random numbers from VEGAS, four phase space evaluations are required. Furthermore, in the limits $x_1 \rightarrow 1$ or $x_2 \rightarrow 1$, there is a cancellation of singularities between the phase space configurations. This leads to an unstable integrand which requires many phase space points to be evaluated correctly.

A second option, first presented in Ref. [98] in the context of CS dipole subtraction, is to instead rescale the PDF functions in the different regions. This can be achieved by introducing a factor $d\xi\delta(\xi - xz)$ for both the z_1 and z_2 integration variables. Schematically the transformation is of the form,

$$\int_0^1 dz f(z, \mu_F^2) \int_0^1 dx M(x, \{xzp_{n+2}\}, \mu_R^2) = \int_0^1 d\xi \int_\xi^1 dx f\left(\frac{\xi}{x}, \mu_F^2\right) M(x, \{\xi p_{n+2}\}, \mu_R^2). \quad (4.2.17)$$

This is a useful transformation because only one phase space evaluation is required for all four of the integration regions. The cost is that the PDFs must now be evaluated four times, however the performance impact of this is relatively small¹. The singularity cancellation in the limits $x_1 \rightarrow 1$ or $x_2 \rightarrow 1$ is now local to a single phase space point, making numerical convergence of the integrand significantly faster. Clearly this method of evaluating the integral is more desirable than rescaling the phase space.

The change in the limits over the integration of the x_1 means that some care must be taken with the plus distributions. Using the second definition of the plus distribution

¹In fact, even if the virtual integration was performed by rescaling the phase space, multiple evaluations of the PDFs might have been required. The momentum fractions would be fixed using this approach, but the factorisation scale μ_F can be dependent on the phase space kinematics. In the case of a dynamic scale choice based on final state observables, the PDFs would have to be evaluated for each integration region.

in Eq. (4.2.10) gives,

$$\int_{\xi}^1 g(x)f(x)_+ = \int_{\xi}^1 dxg(x)f(x) - \int_{\xi}^1 \delta(1-x)g(x)dx \int_0^1 dx'f(x') \quad (4.2.18)$$

$$= \int_{\xi}^1 dxg(x)f(x) - g(1) \int_0^1 dx f(x) \quad (4.2.19)$$

$$= \int_{\xi}^1 dx [(g(x) - g(1))f(x)] - g(1) \int_0^{\xi} dx f(x). \quad (4.2.20)$$

The second term of equation Eq. (4.2.20) can be calculated analytically for all of the plus distributions relevant for the antenna functions,

$$-g(1) \int_0^{\xi} dx \frac{\log^n(1-x)}{1-x} = g(1) \frac{\log^{n+1}(1-\xi)}{n+1}. \quad (4.2.21)$$

In addition to these extra factors that must be included, terms that are proportional to delta functions must also be rescaled,

$$\int_0^1 dxg(x)\delta(1-x) \equiv g(1) = \int_z^1 dx \frac{g(1)}{1-z}. \quad (4.2.22)$$

Using these ingredients, Eq. (4.2.11) can be written as,

$$\begin{aligned} d\sigma^T = & -\mathcal{N}_{\text{NLO}}^V \mathcal{C}(N, n_f) \sum_{ij} \int_0^1 \frac{d\xi_1}{\xi_1} \int_0^1 \frac{d\xi_2}{\xi_2} \frac{1}{S_n} \left\{ \right. \\ & + \int_{\xi_1}^1 dx_1 \int_{\xi_2}^1 dx_2 \left[f_i(\xi_1, \mu_F^2(\{p_n\})) f_j(\xi_2, \mu_F^2(\{p_n\})) \mathcal{X}'_1(x_1, x_2) \right] \\ & + \int_{\xi_1}^1 dx_1 \int_{\xi_2}^1 \frac{dx_2}{x_2} \left[f_i(\xi_1, \mu_F^2(\{p_n\})) f_j \left(\frac{\xi_2}{x_2}, \mu_F^2(\{p_n\}) \right) \mathcal{X}'_2(x_1, x_2) \right] \\ & + \int_{\xi_1}^1 \frac{dx_1}{x_1} \int_{\xi_2}^1 dx_2 \left[f_i \left(\frac{\xi_1}{x_1}, \mu_F^2(\{p_n\}) \right) f_j(\xi_2, \mu_F^2(\{p_n\})) \mathcal{X}'_3(x_1, x_2) \right] \\ & + \int_{\xi_1}^1 \frac{dx_1}{x_1} \int_{\xi_2}^1 \frac{dx_2}{x_2} \left[f_i \left(\frac{\xi_1}{x_1}, \mu_F^2(\{p_n\}) \right) f_j \left(\frac{\xi_2}{x_2}, \mu_F^2(\{p_n\}) \right) \mathcal{X}'_4(x_1, x_2) \right] \\ & \left. \right\} d\Phi_n(p_3, \dots, p_{n+2}; \xi_1 p_1, \xi_2 p_2) M_{n+2}^0(\{\xi_1 \xi_2 p_{n+2}\}) J_n^{(n)}(\{p_n\}), \quad (4.2.23) \end{aligned}$$

where the terms $\mathcal{X}'_1(x_1, x_2), \dots, \mathcal{X}'_4(x_1, x_2)$ are,

$$\begin{aligned} \mathcal{X}'_1(x_1, x_2) = & A'(1, 1) - \sum_i^n \frac{\log^i(1-x_2)}{1-x_2} C'_i(1, 1) - \sum_i^n \frac{\log^i(1-x_2)}{1-x_2} E'_i(1, 1) \\ & + \sum_i^n \sum_j^n \frac{\log^i(1-x_1) \log^j(1-x_2)}{1-x_1} I_{ij}(1, 1), \quad (4.2.24) \end{aligned}$$

$$\begin{aligned} \mathcal{X}'_2(x_1, x_2) &= B'(1, x_2) + \sum_i^n \frac{\log^i(1-x_2)}{1-x_2} C'_i(1, x_2) - \sum_i^n \frac{\log^i(1-x_1)}{1-x_1} G_i(1, x_2) \\ &\quad - \sum_i^n \sum_j^n \frac{\log^i(1-x_1)}{1-x_1} \frac{\log^j(1-x_2)}{1-x_2} I_{ij}(1, x_2), \end{aligned} \quad (4.2.25)$$

$$\begin{aligned} \mathcal{X}'_3(x_1, x_2) &= D'(x_1, 1) + \sum_i^n \frac{\log^i(1-x_1)}{1-x_1} E'_i(x_1, 1) - \sum_i^n \frac{\log^i(1-x_1)}{1-x_1} H_i(x_1, 1) \\ &\quad - \sum_i^n \sum_j^n \frac{\log^i(1-x_1)}{1-x_1} \frac{\log^j(1-x_2)}{1-x_2} I_{ij}(x_1, 1), \end{aligned} \quad (4.2.26)$$

$$\begin{aligned} \mathcal{X}'_4(x_1, x_2) &= F(x_1, x_2) + \sum_i^n \frac{\log^i(1-x_1)}{1-x_1} G_i(x_1, x_2) + \sum_i^n \frac{\log^i(1-x_1)}{1-x_1} H_i(x_1, x_2) \\ &\quad + \sum_i^n \sum_j^n \frac{\log^i(1-x_1)}{1-x_1} \frac{\log^j(1-x_2)}{1-x_2} I_{ij}(x_1, x_2), \end{aligned} \quad (4.2.27)$$

and,

$$\begin{aligned} A'(1, 1) &= \frac{1}{(1-\xi_1)(1-\xi_2)} \left(A(1, 1) + \sum_j^n \frac{\log^{j+1}(1-\xi_2)}{j+1} C_j(1, 1) \right. \\ &\quad \left. + \sum_j^n \frac{\log^{j+1}(1-\xi_1)}{j+1} E_j(1, 1) + \sum_j^n \frac{\log^{j+1}(1-\xi_1)}{j+1} \sum_k^n \frac{\log^{k+1}(1-\xi_2)}{k+1} I_{jk}(1, 1) \right), \end{aligned} \quad (4.2.28)$$

$$B'(1, x_2) = \frac{1}{(1-\xi_1)} \left(B(1, x_2) + \sum_j^n \frac{\log^{j+1}(1-\xi_1)}{j+1} G_j(1, x_2) \right), \quad (4.2.29)$$

$$C'_i(1, x_2) = \frac{1}{(1-\xi_1)} \left(C_i(1, x_2) + \sum_j^n \frac{\log^{j+1}(1-\xi_1)}{j+1} I_{ji}(1, x_2) \right), \quad (4.2.30)$$

$$D'(x_1, 1) = \frac{1}{(1-\xi_2)} \left(D(x_1, 1) + \sum_j^n \frac{\log^{j+1}(1-\xi_2)}{j+1} H_j(x_1, 1) \right), \quad (4.2.31)$$

$$E'_i(x_1, 1) = \frac{1}{(1-\xi_2)} \left(E_i(x_1, 1) + \sum_k^n \frac{\log^{k+1}(1-\xi_2)}{k+1} I_{ik}(x_1, 1) \right). \quad (4.2.32)$$

In the NNLOJET program both of the rescaling methods (rescaling the PDFs and rescaling the phase space) are implemented. Nevertheless, all of the results presented in this thesis were computed by rescaling the PDFs.

4.3 Scale dependence and the running of $\alpha_s(\mu_R^2)$

In section 1.4 it was shown that the entirety of the renormalisation scale dependence for a given observable resides entirely in the scale dependence of $\alpha_s(\mu_R^2)$. The scale dependence of $\alpha_s(\mu_R^2)$ can be inferred from the QCD β -function, evolving from a starting

scale μ_0 to a new scale μ_R . Introducing the definition,

$$L_R = \log \frac{\mu_R^2}{\mu_0^2}, \quad (4.3.1)$$

and substituting this into the QCD β -function leads to,

$$\alpha_s(\mu_0^2) = \alpha_s(\mu_R^2) \left[1 + \beta_0 L_R \left(\frac{\alpha_s(\mu_R^2)}{2\pi} \right) + (\beta_0^2 L_R^2 + \beta_1 L_R) \left(\frac{\alpha_s(\mu_R^2)}{2\pi} \right)^2 + \mathcal{O}(\alpha_s^3(\mu_R^2)) \right]. \quad (4.3.2)$$

Making the α_s dependence of the cross section explicit, the perturbative expansion up to NNLO is given by,

$$\sigma(\mu_0, \alpha_s(\mu_0^2)) = \left(\frac{\alpha_s(\mu_0^2)}{2\pi} \right)^n \sigma^{(0)} + \left(\frac{\alpha_s(\mu_0^2)}{2\pi} \right)^{n+1} \sigma^{(1)} + \left(\frac{\alpha_s(\mu_0^2)}{2\pi} \right)^{n+2} \sigma^{(2)} + \mathcal{O}(\alpha_s^3(\mu_0^2)), \quad (4.3.3)$$

where the individual coefficients $\sigma^{(i)}$, directly analogous to the r^i coefficients in Eq. (1.4.28), have no dependence on μ_0 or $\alpha_s(\mu_0^2)$. n therefore denotes the power of $\alpha_s(\mu_0^2)$ in the leading order contribution to the cross section.

Substituting Eq. (4.3.2) into Eq. (4.3.3) gives,

$$\begin{aligned} \sigma(\mu_R, \alpha_s(\mu_R^2)) &= \left(\frac{\alpha_s(\mu_R^2)}{2\pi} \right)^n \sigma^{(0)} \\ &+ \left(\frac{\alpha_s(\mu_R^2)}{2\pi} \right)^{n+1} \left(\sigma^{(1)} + n\beta_0 L_R \sigma^{(0)} \right) \\ &+ \left(\frac{\alpha_s(\mu_R^2)}{2\pi} \right)^{n+2} \left(\sigma^{(2)} + (n+1)\beta_0 L_R \sigma^{(1)} \right. \\ &\quad \left. + n\beta_1 L_R \sigma^{(0)} + \frac{n(n+1)}{2} \beta_0^2 L_R^2 \sigma^{(0)} \right) \\ &+ \mathcal{O}(\alpha_s^3(\mu_R^2)). \end{aligned} \quad (4.3.4)$$

Hence, for a given perturbative order in $\alpha_s(\mu_R^2)$, the scale dependent terms can be predicted by using only the components from the previous perturbative orders. This proves to be a powerful check on the validity of the calculation, given that the coefficient $\sigma^{(2)}$ in particular contains a large number of scale dependent terms which should cancel.

The scale dependent terms proportional to $\sigma^{(1)}$ occupy different phase spaces because of the real and virtual contributions to the NLO corrections. The scale dependent real contributions reside in the real virtual phase space, whereas the scale dependent virtual contributions lie in the double virtual phase space. There must be a cancellation

between the NLO subtraction terms which resides *within* the full NNLO subtraction terms. This cancellation is non-trivial and forms a powerful check on the consistency of the real virtual and double virtual subtraction terms. It also proves to be a powerful test on the overall colour coefficients. Errors in overall signs or factors would impact adversely on the scale dependence of the final result.

The terms proportional to L_R in Eq. (4.3.4) occupy either the n -parton or $(n+1)$ -parton phase space, as described above. The double real contribution to $\sigma(\mu_R, \alpha_s(\mu_R^2))$ is simply a constant term within $\sigma^{(2)}$. Hence it will be useful to define the quantity,

$$\sigma'^{(2)} = \sigma^{(2)} - \sigma^{(RR)}, \quad (4.3.5)$$

where $\sigma^{(RR)}$ is the contribution from the double real integral with the corresponding $\alpha_s(\mu_R^2)$ dependence removed. We can then define the quantity,

$$\begin{aligned} \sigma'(\mu_R, \alpha_s(\mu_R^2)) &= \left(\frac{\alpha_s(\mu_R^2)}{2\pi}\right)^n \sigma^{(0)} \\ &+ \left(\frac{\alpha_s(\mu_R^2)}{2\pi}\right)^{n+1} \left(\sigma^{(1)} + n\beta_0 L_R \sigma^{(0)}\right) \\ &+ \left(\frac{\alpha_s(\mu_R^2)}{2\pi}\right)^{n+2} \left(\sigma'^{(2)} + (n+1)\beta_0 L_R \sigma^{(1)} \right. \\ &\quad \left. + n\beta_1 L_R \sigma^{(0)} + \frac{n(n+1)}{2} \beta_0^2 L_R^2 \sigma^{(0)}\right) \\ &+ \mathcal{O}\left(\alpha_s^3(\mu_R^2)\right), \end{aligned} \quad (4.3.6)$$

which describes the scale dependence of the real virtual and double virtual contributions to the NNLO calculation. This is a more useful quantity than Eq. (4.3.4) because the double real contribution is the dominant source of statistical uncertainty in the NNLO calculation. Including the contribution from the double real calculation would risk hiding a problem in the scale dependent contribution to the overall calculation.

Fig. 4.2 shows the calculation of Z +jet at NNLO accuracy from the NNLOJET code performed at a fixed scale $\mu_F = M_Z$ and three different renormalisation scales, $\mu_R = [1/2, 1, 2] \cdot M_Z$. This is compared to Eq. (4.3.4) evolving from the fixed scale $\mu_0 = M_Z$. Excellent agreement is observed between the NNLOJET code and the QCD evolution equation.

In principle, the full scale dependence of the calculation can be inferred by also

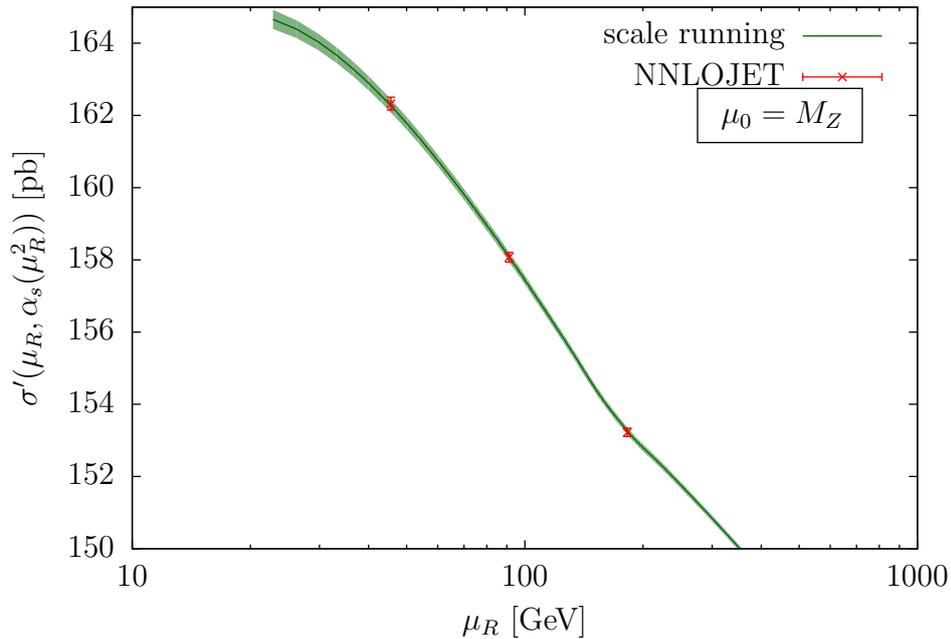


Figure 4.2: The $\sigma'(\mu_R, \alpha_s(\mu_R^2))$ cross section at NNLO accuracy for Z +jet. The green band denotes the prediction from Eq. (4.3.6) running from $\mu_0 = M_Z$, where the area of the band is the statistical uncertainty. The red data points are taken directly from the NNLOJET code ran with $\mu_F = M_Z$ and $\mu_R = [1/2, 1, 2] \cdot M_Z$, with the statistical uncertainty shown.

including contributions from the Altarelli-Parisi evolution equation from a factorisation scale μ_1 to a new factorisation scale μ_F . These terms can be predicted analytically, however in practice extracting the terms from a numerical calculation is extremely challenging up to NNLO. There are many flavour changing contributions to the full evolution equation, making it a lengthy exercise to validate these terms.

4.4 Double real numerical integration

The double real contribution to the NNLO calculation is the most challenging to construct and evaluate numerically. It contains the highest dimensionality phase space in the problem and also requires a sufficient number of evaluations to resolve the entire double unresolved phase space. In terms of statistical uncertainties it is usually far larger than either the real virtual or double virtual integrals, hence it should be the focus of any attempt to numerically optimise the calculation.

Tab. 4.1 demonstrates an approximate breakdown of the CPU time required for a

| Integral | Approximate CPU time [hours] |
|------------------------|------------------------------|
| Double real region a | 75,000 |
| Double real region b | 75,000 |
| Real virtual | 10,000 |
| Double virtual | 500 |
| Real | 1,000 |
| Virtual | 100 |
| Leading order | < 1 |
| Total | 161,601 |

Table 4.1: Approximate CPU hours each of the contributions to the full NNLO calculation of Z +jet. The double real phase space is divided into two regions, a and b , to optimise the sampling of distinct double unresolved configurations.

full NNLO calculation of Z +jet production. The double real phase space is divided into two regions, a and b , to optimise the sampling of distinct double unresolved configurations. Further details on the division of the phase space can be found in section 4.4.1.

The double real phase space calculation is a two stage process,

- $\mathcal{O}(10^7)$ events are used to warmup the VEGAS grid such that the sampling of the integrand is optimised as far as possible.
- The VEGAS grid is frozen and is used to evaluate $\mathcal{O}(10^8)$ events in a production run which are used to evaluate the total cross section and distributions.

Given the slow evaluation time of the integrand, it is not feasible to run the code on a single core for either the warmup or the production runs.

A warmup run consists of many iterations, typically 10, where between each iteration the weights from the previous iterations are used to adapt the sampling, as according to the VEGAS algorithm [149]. This means that to perform a multithreaded warmup run the threads must at least communicate after each iteration. Two technologies are readily available to evaluate a multithreaded integral,

- openMPI², where messages are passed between a main thread and the subsequent daughter threads. The memory for each daughter thread is cloned in each case

²<http://www.mcs.anl.gov/research/projects/mpi/>

and requires a special compiler with specific code modifications to run. It is limited by the number of cores available in a cluster.

- OpenMP³, where the memory is shared amongst the threads and can be run using the conventional GNU compiler. OpenMP is limited by the number of cores available in a single node.

A hybrid code, using both openMPI and openMP, is also possible. For the purposes of evaluating $\mathcal{O}(10^7)$ phase space points to generate the VEGAS grid openMP is sufficient. Using 16 cores, the evaluation time of a warmup run is approximately 20 hours for each phase space region.

For a production run, using a frozen VEGAS grid, no communication is required between the threads. So that we are not limited by the use of a single cluster and can instead use the full resources available via the UK grid contribution [150], we avoid using openMP and openMPI for production runs. Instead the number of events are divided into N independent runs, each evaluating M events. For a typical production run, the number of runs N is approximately $\mathcal{O}(10^3)$ and M is $\mathcal{O}(10^5)$ events. Each run takes approximately 10 – 20 hours depending on the performance of the cores in question⁴. The method of combining so many runs together is the topic of discussion for chapter 5.

4.4.1 Double real phase space regions

As discussed in the previous section, the double real phase space is divided into two regions, a and b , to evaluate different double unresolved configurations as efficiently as possible.

³<http://openmp.org>

⁴The evaluation of the integrand is significantly slower during production than compared to warmup runs. This is because producing histograms introduces a huge overhead into the code.

Region a

Region a focuses primarily on generating triple collinear limits. The generation is of the form,

$$\begin{aligned} p_{\hat{1}} + p_{\hat{2}} &\rightarrow p_{345} + p_{\ell\ell} \\ &\rightarrow p_3 + p_{45} + p_{\ell^+} + p_{\ell^-} \\ &\rightarrow p_3 + p_4 + p_5 + p_{\ell^+} + p_{\ell^-}. \end{aligned}$$

This phase space can then be isolated into wedges by placing kinematic cuts on the invariants of the phase space. This is useful because we can isolate the divergent limits within a specific wedge and can subsequently generate the appropriate rotated phase spaces to cancel the azimuthal terms.

The primary wedge of region a is designed to isolate the triple collinear limits $3||4||5$ and $\hat{1}||4||5$; and also the 3 soft, $4||5$ collinear limits. This is done by applying constraints on the smallest invariant and second smallest invariant, s_1 and s_2 respectively,

$$(\{s_1, s_2\} \in \{s_{45}, s_{\hat{1}4}, s_{\hat{1}5}\}),$$

or,

$$(s_1 = s_{45}, s_2 = \min(s_{34}, s_{35}) \text{ and } s_{\hat{1}4} + s_{\hat{1}5} < s_{24} + s_{25}).$$

The full phase space can then be recovered by permuting parton 3 through the final state⁵ and exchanging the initial state partons $\hat{1}$ and $\hat{2}$. Therefore, in total, there are 6 wedges to evaluate for a given event.

Since we are isolating the $3||4||5$ and $\hat{1}||4||5$ triple collinear limits, the appropriate rotations to perform are rotations around the shared triple collinear $3||4||5$ and $\hat{1}||4||5$ axes, as was suggested with the spike testing discussed in section 4.1. Unfortunately, rotations such that those are challenging to implement within a general phase space generator. Most rotations will change the phase space weight because the invariants are changed within the event. These changes can be difficult to manage, particularly with processes involving a boost from rotating an initial state parton. What can be

⁵The wedge cuts are symmetric on partons 4 and 5 and therefore a full permutation of the final state is not necessary.

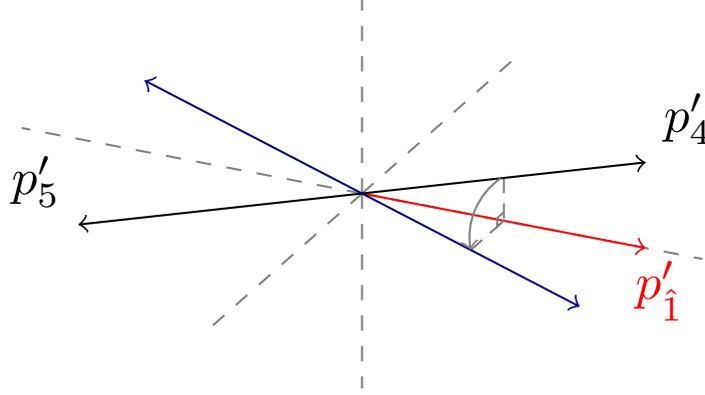


Figure 4.3: The rotation by $\pi/2$ applied in the phase space generator around the direction of the initial state parton p'_1 in the rest frame of p_{45} .

done however is a rotation around a shared axis within the rest frame of a decaying particle. Consider the decaying frame of the p_{45} particle, as shown in Fig. 4.3. The initial state particle $p_{\hat{1}}$ can be boosted into this frame to $p'_{\hat{1}}$ and is defined to lie on the z axis in this frame. A rotation of $\pi/2$ is then performed around the z axis. In so doing, the invariants $s_{\hat{1}4}$, $s_{\hat{1}5}$ and s_{45} are left unchanged and therefore the weight of the phase space remains the same.

Given that these rotations are not the perfect choice, it is important to check that they do successfully cancel the azimuthal terms. Fig. 4.4(b) shows a spike plot for a triple collinear initial-final-final $\hat{1}_g||3_g||4_g$ limit using the rotation depicted in Fig. 4.3. This can be compared to Fig. 4.4(c), which is the same limit with the same unrotated events using the rotation described in section 4.1. The performance of the subtraction term is better in Fig. 4.4(c) compared to Fig. 4.4(b). Nevertheless, in comparison to not applying any rotations (shown in Fig 4.4(a)) there are huge gains in the convergence of the subtraction term.

To cancel azimuthal terms related to the $\hat{1}||4||5$ limit in region a , 4 phase spaces must be evaluated per wedge,

- the unrotated phase space,
- p_{345} rotated about $p_{\hat{1}}$,
- p_{45} rotated about $p_{\hat{1}}$,
- p_{345} rotated about $p_{\hat{1}}$ and p_{45} rotated about $p_{\hat{1}}$,

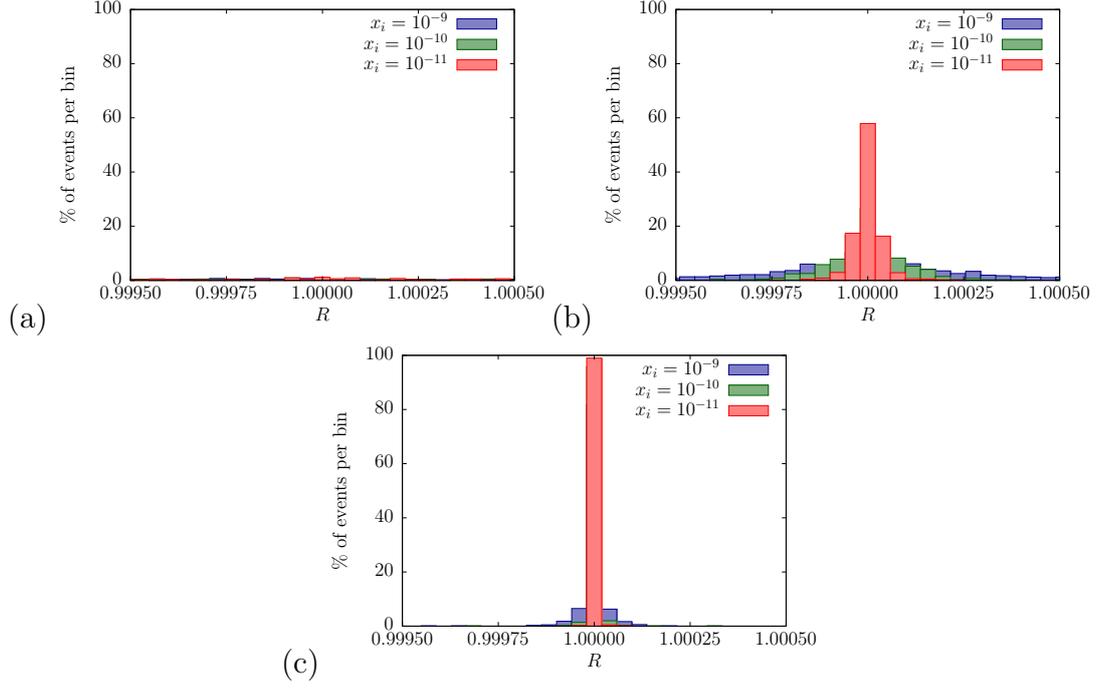


Figure 4.4: Spike plots for a triple collinear initial-final-final $\hat{1}_g||3_g||4_g$ limit for the $qg \rightarrow qggZ$ leading colour contribution to Z +jet production with (a) no rotations applied, (b) a rotation about the p_1 axis and (c) a rotation about the shared $p_1 + p_3 + p_4$ collinear axis, boosted back on beam. The phase space points generated before the rotation is applied are identical in the three spike plots.

therefore in total there are 24 phase space evaluations to be made per event. Note that no rotations are applied to cancel the azimuthal correlations in the $3||4||5$ limit. In principle, one could apply further rotations to regulate $3||4||5$ limit, however the numerical cost to evaluate many more phase space points is significant. Indeed, it is more efficient to not cancel the azimuthal correlation terms through rotations for the $3||4||5$ limit and instead let them cancel non-locally.

Region b

Region b focuses on generating double collinear limits. The generation of the phase space points is precisely the same as region a , however the wedge cuts are different. The selection for the primary wedge criteria are,

$$(s_1 = s_{13} \text{ and } s_2 = \min(s_{23}, s_{24}, s_{25}))$$

or,

$$(\{s_1, s_2\} \in \{s_{45}, s_{13}\}),$$

which selects the double collinear limits $\hat{1}||3$ and $4||5$, $\hat{1}||3$ and $\hat{2}||4$, $\hat{1}||3$ and $\hat{2}||5$. Similarly to region a , the full phase space is recovered from a permutation of parton 3 through the final state and the exchange of the initial state partons $\hat{1}$ and $\hat{2}$.

Within this wedge of region b , the rotations that are applied are a rotation of the p_{345} system about the p_1 axis and a rotation of the p_{45} system about the p_2 axis. Similarly to region a , this implies that there are 4 phase spaces to be evaluated per wedge,

- the unrotated phase space,
- p_{345} rotated about p_1 ,
- p_{45} rotated about p_2 ,
- p_{345} rotated about p_1 and p_{45} rotated about p_2 ,

and hence there are in total 24 phase space points to be evaluated per event.

4.4.2 Double real phase space sampling

The division of the double real phase space into two regions is also important so that very distinct configurations of the phase space can be sampled in their entirety. The iterated decay procedure might be identical in region a and region b , but it would be impossible to sample phase space configurations that favourably satisfy both sets of wedge cuts simultaneously. From this perspective the two regions represent a simplified multichannel integration procedure [151], where we can perform a change of variables in each case to sample high weight contributions to the primary wedges as much as possible.

Each new decay in the phase space generates three new variables to sample, (s, θ, ϕ) , which are defined in the rest frame of a decaying particle of mass \sqrt{s} . The angle θ can be redefined in terms of other invariants in the phase space. These invariants can then be sampled logarithmically so that the dominant contributions to the integral are sampled more frequently.

p_{345} decay sampling

Consider first the phase space generated from the decay,

$$p_{345} \rightarrow p_3 + p_{45},$$

where p_{45} is off-shell with mass $\sqrt{s_{45}}$. As discussed previously, the initial state parton $p_{\hat{1}}$ is boosted into the rest frame of p_{345} such that it is defined to be on the z axis. Therefore by definition,

$$p'_{\hat{1}} = E'_1(1, 0, 0, 1), \quad (4.4.1)$$

where E'_1 is the energy of parton 1 in the rest frame of p_{345} . It follows that,

$$p'_3 = E'_3(1, \sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta), \quad (4.4.2)$$

$$p'_{45} = (E'_{45}, -E'_3 \sin \theta \cos \phi, -E'_3 \sin \theta \sin \phi, -E'_3 \cos \theta), \quad (4.4.3)$$

where,

$$E'_3 = \frac{s_{345} - s_{45}}{2\sqrt{s_{345}}}, \quad (4.4.4)$$

$$E'_{45} = \frac{s_{345} + s_{45}}{2\sqrt{s_{345}}}. \quad (4.4.5)$$

From these definitions, it is clear that,

$$2p'_{\hat{1}} \cdot p'_3 \equiv -s_{\hat{1}3} = 2E'_1 E'_3 (1 - \cos \theta), \quad (4.4.6)$$

$$2p'_{\hat{1}} \cdot p'_{45} \equiv -(s_{\hat{1}4} + s_{\hat{1}5}) = 2E'_1 (E'_{45} + E'_3 \cos \theta). \quad (4.4.7)$$

Combining Eqs. (4.4.6) and (4.4.7) gives,

$$\cos \theta = 1 + \frac{s_{\hat{1}3}}{2E'_1 E'_3} \equiv 1 - \frac{2s_{\hat{1}3}}{s_{\hat{1}3}|_{\max}}, \quad (4.4.8)$$

which defines $s_{\hat{1}3}|_{\max}$ as the kinematic maximum of $s_{\hat{1}3}$. In terms of invariants this is,

$$\cos \theta = 1 - \frac{2s_{\hat{1}3}s_{345}}{(s_{345} - s_{45})(s_{\hat{1}345} - s_{345})}, \quad (4.4.9)$$

such that,

$$s_{\hat{1}3}|_{\max} = \frac{(s_{345} - s_{45})(s_{\hat{1}345} - s_{345})}{s_{345}}. \quad (4.4.10)$$

In region b , Eq. (4.4.8) is used to sample $s_{\hat{1}3}$ logarithmically, including an appropriate Jacobian factor to make the change of variables from $d\theta$ to $ds_{\hat{1}3}$,

$$d \cos \theta = 2 \frac{ds_{\hat{1}3}}{s_{\hat{1}3}|_{\max}}. \quad (4.4.11)$$

Clearly to sample $s_{\hat{1}3}$ logarithmically requires some kinematical cutoffs to be applied. $s_{\hat{1}3}$ is bounded from above by $s_{\hat{1}3}|_{\max}$, however from below $s_{\hat{1}3}$ can go arbitrarily small. We must therefore introduce a non-zero minimum to $s_{\hat{1}3}$,

$$s_{\hat{1}3}|_{\min} = y_0 \cdot \hat{s}, \quad (4.4.12)$$

where y_0 is a small dimensionless parameter. This is a technical cutoff on the phase space such that phase space configurations with $s_{\hat{1}3} < y_0 \cdot \hat{s}$ are no longer generated. This can be motivated by the fact that as $s_{\hat{1}3} \rightarrow 0$, the weight of the event becomes negligible since there is a large cancellation between the matrix element and subtraction term. More details on the technical cut dependence can be found in section 4.5.

In region a , a further transformation is applied. The variable r , defined as,

$$r = \frac{s_{\hat{1}3}}{s_{\hat{1}3}|_{\max} - s_{\hat{1}3}}, \quad (4.4.13)$$

is sampled logarithmically instead of $s_{\hat{1}3}$ directly. The denominator of r can be rewritten in terms of other invariants,

$$s_{\hat{1}3}|_{\max} - s_{\hat{1}3} = \left(1 + \frac{s_{45}}{s_{345}}\right) (s_{\hat{1}4} + s_{\hat{1}5}) - \frac{s_{45}}{s_{345}} (s_{\hat{1}345} - s_{345}), \quad (4.4.14)$$

where the free parameter of Eq. (4.4.14) is $s_{\hat{1}4} + s_{\hat{1}5}$. This means that if we sample r logarithmically, we sample configurations where either $s_{\hat{1}3}$ is small (appropriate for the 3 soft, 4||5 collinear limit) *or* where $s_{\hat{1}4} + s_{\hat{1}5}$ is small (appropriate for the $\hat{1}$ ||4||5 triple collinear limit).

The problem with the definition of Eq. (4.4.14) is the kinematic boundary where $s_{\hat{1}3} = s_{\hat{1}3}|_{\max}$. Here the ratio r is ill-defined. Furthermore, $s_{\hat{1}3}|_{\max} - s_{\hat{1}3}$ has a non-trivial dependency on y_0 as invariants approach the technical cutoff. This means that it is hard to introduce a bound on $s_{\hat{1}3}|_{\max} - s_{\hat{1}3}$ that does not artificially restrict the sampling of the phase space. A pragmatic solution to the problem is to modify r such

that,

$$r = \frac{s_{\hat{1}3}}{s_{\hat{1}3}|_{\max} - s_{\hat{1}3} + y_0 \hat{s}}, \quad (4.4.15)$$

which means that $s_{\hat{1}3}$ can be exactly equal to the kinematic boundary of $s_{\hat{1}3}|_{\max}$. This definition will sample favourably where $s_{\hat{1}4} + s_{\hat{1}5}$ is close to $y_0 \hat{s}$.

Using these definitions, the Jacobian factor for changing $ds_{\hat{1}3}$ to dr is,

$$ds_{\hat{1}3} = \frac{(s_{\hat{1}3}|_{\max} - s_{\hat{1}3} + y_0 \hat{s})^2}{s_{\hat{1}3}|_{\max} + y_0 \hat{s}} dr. \quad (4.4.16)$$

p_{45} decay sampling

The subsequent decay of,

$$p_{45} \rightarrow p_4 + p_5,$$

is a simpler scenario than the p_{345} decay considered previously since p_4 and p_5 are both on-shell. For region a , $p_{\hat{1}}$ is defined to be on the z axis and therefore the angle θ is simply given by,

$$\cos \theta = 1 - \frac{2s_{\hat{1}4}}{s_{\hat{1}4} + s_{\hat{1}5}} \equiv 1 - \frac{2s_{\hat{1}4}}{s_{\hat{1}4}|_{\max}}, \quad (4.4.17)$$

where $s_{\hat{1}4}|_{\max}$ is $s_{\hat{1}4} + s_{\hat{1}5}$ and is an input variable from the p_{345} decay. Hence for region a the appropriate variable to sample logarithmically is,

$$r = \frac{s_{\hat{1}4}}{s_{\hat{1}5}}, \quad (4.4.18)$$

making the appropriate change of variables to change from $d \cos \theta$ to dr ,

$$d \cos \theta = 2 \frac{ds_{\hat{1}4}}{s_{\hat{1}4}|_{\max}}, \quad (4.4.19)$$

$$= 2 \frac{s_{\hat{1}5}^2 dr}{s_{\hat{1}4} s_{\hat{1}4}|_{\max}}. \quad (4.4.20)$$

This will sample configurations where either $s_{\hat{1}4}$ or $s_{\hat{1}5}$ are small. The maximum and minimum boundaries on the variable r are defined as,

$$r|_{\min} = \frac{y_0 \hat{s}}{s_{\hat{1}4}|_{\max}}, \quad (4.4.21)$$

$$r|_{\max} = \frac{s_{\hat{1}4}|_{\max}}{y_0 \hat{s}}. \quad (4.4.22)$$

The change of variables for region b is very similar. Unlike in region a , $p_{\hat{2}}$ is now

defined to be on the z axis. Hence it follows that,

$$\cos \theta = 1 - \frac{2s_{\hat{2}4}}{s_{\hat{2}4} + s_{\hat{2}5}} \equiv 1 - \frac{2s_{\hat{2}4}}{s_{\hat{2}4}|_{\max}}, \quad (4.4.23)$$

and therefore to sample both $s_{\hat{2}4}$ and $s_{\hat{2}5}$ favourably, we can define r such that,

$$r = \frac{s_{\hat{2}4}}{s_{\hat{2}5}}, \quad (4.4.24)$$

where the Jacobian factors and boundaries on r are precisely the same as for region a with parton $\hat{1}$ exchanged for parton $\hat{2}$.

4.5 Technical cut dependence of real integrals

Analytically the subtraction term is constructed to regulate phase space singularities in the absolute unresolved limit. Numerically it is not feasible to probe absolute unresolved limits because numbers are only known to within n digits to the processor at runtime. In a numerical evaluation of the form,

$$A - A + \delta, \quad (4.5.1)$$

where $A > \delta \cdot 10^n$, there will be a loss of precision in the final evaluation of this expression. These sort of calculations should be avoided as much as possible in numerical calculations.

The NNLOJET code is written using double precision variables in FORTRAN, meaning that $n \sim 16$. It is possible to use quadruple precision variables to increase this to $n \sim 34$, however this would have a negative impact on the performance of the code. A much more desirable solution to the problem is to prevent the code from evaluating extremely close to singular regions of the phase space where the numerical cancellation between the matrix element and subtraction term are largest. The weight of the event in the absolute unresolved limit must be zero, therefore we can apply a technical cutoff of the integrand a small but finite distance away from the singular region. This will automatically set the integrand to zero when the technical cutoff constraint is satisfied.

The method by which the technical cutoff is applied is a choice. In the NNLOJET

code, the choice of technical cut constraint is given by,

$$s_{\min} < y_0 \cdot \hat{s}, \quad (4.5.2)$$

where s_{\min} is the smallest invariant in the full phase space, y_0 is a small dimensionless parameter and \hat{s} is the partonic centre-of-mass energy. If this condition is satisfied then the entire event is vetoed.

Using this technical cut definition, y_0 is a free parameter. In terms of numerical stability, y_0 should be set as high as possible to avoid probing low weight regions of the phase space with a large cancellation between the matrix element and subtraction term. The problem is that if y_0 is set too high then non-negligible contributions to the cross section will be cut away. A balance must be struck between a statistical uncertainty of probing unresolved configurations and a systematic uncertainty with a technical cut dependence.

Unfortunately, a ‘good’ choice for y_0 cannot be inferred *a priori*, it can only be determined by performing multiple evaluations of the integral in question with different values for y_0 . If the answer is independent of the choice of y_0 (within statistical errors) then the subtraction term is performing well and the technical cut has been set to a reasonable value.

In principle, it is always possible to resolve a technical cut dependence using an infinite number of evaluations of the integral. Nevertheless, if it has no impact on the physical conclusions to be drawn from the calculation then it is irrelevant. For example, if there is a systematic shift of 0.1% in the calculation due to the technical cut and this is compared to data with 1% uncertainties then the shift due to the technical cut can be safely ignored. Clearly this is highly dependent on the calculation in question and what its intended application is. A great deal of forethought must go into the choice of technical cut.

4.5.1 Reduced phase space technical cut

As discussed in section 2.4.1, there are many terms within the double real subtraction term which consist of an iterated $X_3^0 \otimes X_3^0$ antenna structure. The final state phase

space mappings for these terms are of the form,

$$\{p_{n+2}\} \rightarrow \{p_{n+1}\} \rightarrow \{p_n\}, \quad (4.5.3)$$

where the observables for this term are evaluated from the n -parton phase space.

From this iterated mapping structure, it is always possible to generate a single unresolved configuration from a well resolved $(n+2)$ -parton phase space. As an example, consider the initial-final mapping function $(\hat{1}, 3, 4) \rightarrow (\hat{1}, \widetilde{(34)})$ applied to the $(n+2)$ -parton phase space. The momentum of the final state composite parton $\widetilde{(34)}$ is given by,

$$p_{\widetilde{(34)}} = p_3^\mu + p_4^\mu - (1-x)p_1^\mu, \quad (4.5.4)$$

where x is defined in Eq. (2.1.14). The invariant between this composite parton and the second initial state parton $\hat{2}$ is therefore,

$$s_{\hat{2}\widetilde{(34)}} = s_{\hat{2}3} + s_{\hat{2}4} - (1-x)\hat{s}. \quad (4.5.5)$$

This invariant is small when $s_{\hat{2}3} + s_{\hat{2}4} \rightarrow (1-x)\hat{s}$. For this condition to be satisfied, none of the invariants in the $(n+2)$ -partons phase space need to be small. Similar conclusions can be drawn from the initial-initial and final-final mapping functions.

These configurations pose a serious problem for the convergence of the double real integrand. The double real matrix element will be smooth and well behaved in this region, whereas the subtraction term will be generating artificially unresolved phase space points which, due to the non-locality of the internal cancellation of the antenna functions, can generate a large unphysical weight.

Since these spurious divergent configurations must cancel within the subtraction term, an appropriate solution is to apply a technical cut to the reduced $(n+1)$ -parton phase space in a similar fashion to the previous section. Some care must be made in the choice of this technical cut however. If, for example, we choose a reduced technical cut that matches the technical cut of the full phase space ($s_{\min} < y_0 \cdot \hat{s}$) then we introduce numerical instabilities on the double unresolved phase space boundary. The natural cancellation between the terms on the double unresolved boundary would be spoiled as terms are switched on and off as phase spaces are mapped above and below the technical cut. This means a tighter cut is required on the $(n+1)$ -parton phase space

compared to the full phase space. As a first attempt at a technical cut, a reasonable cut to introduce is,

$$s'_{\min} < \frac{y_0}{100} \hat{s}, \quad (4.5.6)$$

where s'_{\min} occupies the $(n+1)$ -parton phase space. Having both the full and reduced technical cuts related to the same y_0 is useful because varying y_0 will check the dependency of the integral on both cuts. When an $(n+1)$ -parton phase space fails the reduced technical cut, only that contribution to the subtraction term is vetoed. The full matrix element and all of the other terms within the subtraction term are evaluated as per usual.

An extension of this is to consider a dynamically varying cut on the $(n+1)$ -parton phase space. The motivation for introducing a reduced technical cut is to remove mapped phase space configurations that are more divergent than the full $(n+2)$ -parton phase space. This is a relative statement between the full and reduced phase space, hence a more appropriate technical cut is a measure of the unresolved configuration relative to the full phase space. To that end, we can construct a reduced technical cut of the form,

$$s'_{\min} < \max \left(\delta_0 \cdot s_{\min}, \frac{y_0}{100} \hat{s} \right), \quad (4.5.7)$$

where s_{\min} is the smallest invariant in the $(n+2)$ -parton phase space and δ_0 is a small parameter. This accommodates the need to remove phase spaces that are relatively unresolved in comparison to the $(n+2)$ -parton phase space.

It should be stressed that when the reduced technical cut is triggered, the full matrix element is still evaluated. The reduced technical cut should not negatively impact the convergence of the integral. Nevertheless, numerical checks must be made to find an appropriate value for δ_0 which does not introduce a systematic dependence on this cut.

Fig. 4.5 demonstrates the total cross section calculation for Z +jet double real contributions for different values of y_0 , where the reduced technical cut has been fixed such that $\delta_0 = 0$ (i.e. just the reduced technical cut of Eq. (4.5.6) is applied). In both region a and b the calculation is stable as a function of y_0 , suggesting that the subtraction term is performing well and that technical cuts around these values are a sensible choice. For the results presented in chapters 5 and 6, y_0 is fixed to be 10^{-6} .

Fig. 4.6 shows the contribution to the cross section for different values of δ_0 . The

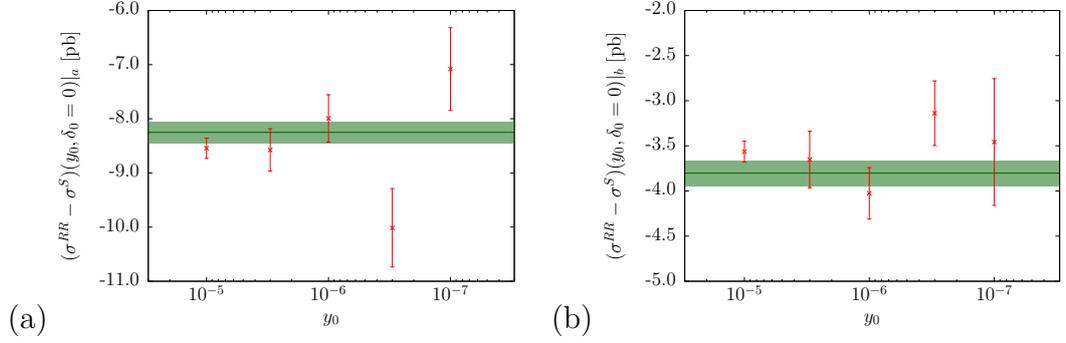


Figure 4.5: Contributions to the cross section for different values of the full technical cut, y_0 , for (a) double real region a and (b) double real region b , where the error bars on the datapoints are statistical. Approximately the same number of events were used for each datapoint ($\mathcal{O}(10^8)$). The green line represents an equivalent calculation with significantly more events ($\mathcal{O}(10^9)$) for $y_0 = 10^{-6}$, where the green band denotes its corresponding statistical uncertainty.

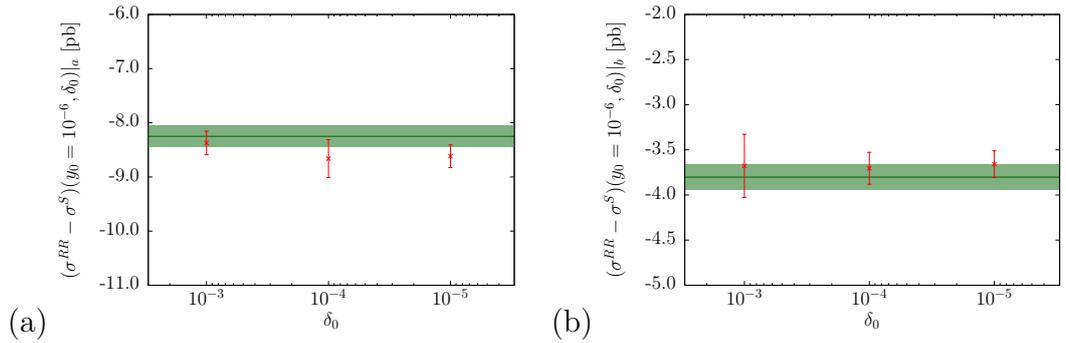


Figure 4.6: Contributions to the cross section for different values of δ_0 for (a) double real region a and (b) double real region b , where the error bars on the datapoints are statistical. Approximately the same number of events were used for each datapoint ($\mathcal{O}(10^8)$) and the full technical cut was set to $y_0 = 10^{-6}$. The green line represents an equivalent calculation with significantly more events ($\mathcal{O}(10^9)$) for $y_0 = 10^{-6}$ and $\delta_0 = 0$, where the green band denotes its corresponding statistical uncertainty.

full technical cut has been fixed as $y_0 = 10^{-6}$ for all of the datapoints. For both region a and b , Fig. 4.6 shows an excellent agreement between the results for various choices of δ_0 and the high statistics result with $\delta_0 = 0$. Note that the datapoints in Fig. 4.6 are of a comparable size to the run using approximately an order of magnitude more events. This shows that the dominant source of statistical uncertainty from the double real integration were from terms where the $(n+1)$ -parton phase space was more unresolved than the full phase space.

Unfortunately the results in chapters 5 and 6 were derived before the δ_0 improvement to the reduced technical cut was implemented. Nevertheless this is a promising significant upgrade for future runs of the Z +jet calculation which shows no remaining dependence on δ_0 .

4.6 NNLO automation techniques

Even for the simplest processes, implementation of the antenna formalism poses a significant bookkeeping challenge. Each initial state combination and colour contribution must be independently tested and subsequently combined to construct the cross section up to NNLO accuracy. A great deal of forethought is required to make the numerical implementation as simple as possible.

In principle, it is highly desirable to automate the construction of numerical routines as far as is feasibly possible. As discussed previously, both the CS and FKS subtraction methods have been fully automated for NLO subtraction [98, 99]. At NNLO, the overlapping singularity structure makes a fully automated approach far more challenging and is not really feasible in the short term. Nevertheless, there are extremely useful auto-generation tools that can be used to assist in the management and construction of the subtraction terms. This section will focus on two distinct auto-generation tools used to make a full NNLO implementation simpler,

- auto-generation of FORTRAN, L^AT_EX and FORM versions of a subtraction term from a common scripting language,
- auto-generation of driver routines to manage the basic structure of the integration code.

4.6.1 Auto-generation of subtraction terms

For a non-trivial process such as Z +jet production, it is common for NNLO subtraction terms to have $\mathcal{O}(100)$ terms for a single colour contribution and initial state configuration. In FORTRAN code, each contribution to the subtraction term is of the form,

```
* {i4 = j3 , i5 = j4 , i6 = j5 , i7 = j6 , [i1] = j1 , [i2] = j2}
   call pmap7to6II(i1 , i2 , i3 , i4 , i5 , i6 , i7 , j1 , j2 , j3 , j4 , j5 , j6 , ipass)
* {j3 = k3 , j5 = k4 , j6 = k5 , [j1] = k1 , [j2] = k2}
   call pmap6to5II(j1 , j2 , j4 , j3 , j5 , j6 , k1 , k2 , k3 , k4 , k5 , ipass)
   call ecuts_vj(5 , ipass)
   if(ipass.eq.1)then
     jpass(18)=1
     call getqcdnorm(ix , partons(18 , :) , facnorm(18 , :))
     wt(18)=+1d0*FullqgD30II(i1 , i3 , i2 , 7) * FullqgA30II(j1 , j2 , j4 , 6)
       *B1g0Z(k1 , k2 , k3 , k4 , k5)
     call bino(1 , partons(18 , :) , -relfac*wt(18)*facnorm(18 , :) , 5)
   endif
```

and therefore there is a high probability of introducing typographical errors. Given the large number of possible initial state configurations and colour levels, it quickly becomes impossible to write all of the FORTRAN code required for a numerical implementation by hand.

A better solution is to write the subtraction term using a short scripting language, which in turn gets parsed into FORTRAN code. This can be achieved using a set of scripts written in Maple⁶, which convert a line of the form,

```
+qgD30II(1 , i , 2) * qgA30II([1] , [2] , Q) * B1g0Z([[1]] , [[2]] , j , Z) * JET11(j) *
  a18
```

into the FORTRAN code shown above. Each subtraction term is written to a separate ‘map’ file which can be outputted as FORTRAN code. This allows a subtraction term to be debugged and implemented quickly. Furthermore, any change in the syntax of the FORTRAN code (e.g. to add a new feature) requires only a single modification of

⁶<https://www.maplesoft.com/>

the controlling Maple script and *all* of the subtraction terms in the NNLOJET code are updated.

Another useful feature that arises from this approach is that the same scripts can be used to output L^AT_EX code. This is extremely useful because it guarantees that the subtraction term that is documented is *precisely* the same as is being used within the NNLOJET code. The above term is generated as,

$$+ D_{3,gg}^0(1, i, 2) A_{3,gg \rightarrow qq}^0(\bar{1}, \bar{2}, Q) B_1^{Z,0}(\bar{1}, \bar{2}, j, Z) J_1^{(1)}(\{p\}_1)$$

For processes involving ϵ pole structures from loop integration, it is also useful to output FORM code so that the cancellation of the ϵ poles between the matrix element and the subtraction term can be checked analytically. The Maple scripts described above will also output FORM code for these processes which can be used to check the explicit pole cancellation.

A fully automated approach to antenna subtraction would be if the inputted ‘.map’ files were generated automatically. If the antenna functions had precisely the desired limits then this would be a very simple exercise. Unfortunately the complicated nature of the spurious singularity structure of certain antenna functions (e.g. the D_4^0 antenna discussed in section 2.5) makes this extremely challenging.

4.6.2 Auto-generation of driver routines

The NNLOJET code is constructed such that the numerical evaluations of the subtraction term and the matrix element are factorised in terms of PDF contributions and colour levels. Each combination of colour level and initial state configuration is considered to be a separate channel which we sum over to generate the full NNLO result. The reason for this choice is partially out of necessity with the construction of the subtraction terms, but it is also useful to keep the code as flexible as possible. A user can request specific contributions, for example to make a leading colour approximation or a specific initial state configuration, and can calculate only these terms. It will be shown later in section 4.7.2 how this flexibility can be exploited to gain a huge performance increase.

The cost however in having this degree of flexibility is that there are a vast number

of channels that need to be introduced into the code. This is a largely repetitive structure which can be auto-generated with ease. Indeed, all that needs to be inputted is the name of a given matrix element, its parton content and its relative colour factor $\mathcal{C}(N, n_f)$ and a series of Maple scripts will generate all of the appropriate initial state channels. For example, a list of the form,

```
[sB3g0Z , [ qb , g , g , g , q , lp , lm ] , 1 ] ,
```

will generate *all* of the channels for this particular colour contribution, taking into account averaging factors over identical final state partons. The first element of this list is the name of the matrix element which follows the naming convention which will be outlined in section 4.6.3. The second element is the particle content (in this example $q, g, g, g, \bar{q}, \ell^+, \ell^-$) and the final term is the relative colour factor $\mathcal{C}(N, n_f)$. This particular example is a leading colour contribution, hence $\mathcal{C}(N, n_f) = 1$.

In the outputted FORTRAN code, each call to the matrix element for a specific channel is of the form,

```
c—— d db to g g g lp lm
      if(ip(163).and.coloursel(1))then
        iproc = 163
        nfB1 = 1
        ip1 = 10
        ip2 = -10
        call getqcdnorm(ix , partons , factor)
        kinwt = factor*(sB3g0Z(1,3,4,5,2,6,7)
        .           +sB3g0Z(1,3,5,4,2,6,7)
        .           +sB3g0Z(1,4,3,5,2,6,7)
        .           +sB3g0Z(1,4,5,3,2,6,7)
        .           +sB3g0Z(1,5,3,4,2,6,7)
        .           +sB3g0Z(1,5,4,3,2,6,7))
        wt = partons*kinwt
        call bino(1, partons , kinwt , 5)
        call storewgt(1, wt)
        wtsum = wtsum + wt
      endif
```

which consists of a sum over all of the colour orderings within a specific channel. This particular channel is the leading colour contribution $d\bar{d} \rightarrow ggg$ to Z +jet at NNLO. The corresponding subtraction term would be,

```

if (ip(163) .and. colourSel(1)) then
  iproc = 163
  nfB1 = 1
  ip1 = 10
  ip2 = -10
  wt = qqbsB3g0ZS(1,3,4,5,2,6,7)
  wtsum = wtsum + wt
endif

```

where the routine,

```
qqbsB3g0ZS(1,3,4,5,2,6,7)
```

in this channel is a blank wrapper routine for the full subtraction term. Routines generated using the scripts described in section 4.6.1 can then be easily imported into this blank routine. This is the only part of the NNLOJET code that requires user intervention when adding new processes.

In addition to these two routines, there is a third auto-generated subroutine called ‘getqcdnorm’ that processes the overall factor of the particular channel. In this example there are three identical gluons in the final state, hence the code generates an additional factor of $1/3!$,

```

case (163)
  factor=1d0/6d0*facRR ! d db -> g g g lp lm sB3g0Z

```

The workflow of construction terms for the NNLOJET program is summarised in Fig. 4.7. When a channel needs to be calculated, a corresponding ‘.map’ file is constructed to regulate the singularities of the given channel. The ‘.map’ file is used to generate FORTRAN routines which can be tested immediately by generating spike plots and/or checking the ϵ pole cancellation. This forms a feedback loop, where whether the subtraction term is working or not is inferred by the quality of the spike plots produced and whether the explicit ϵ poles cancel.

Once the spike plots look good enough in all of the singular regions of the phase

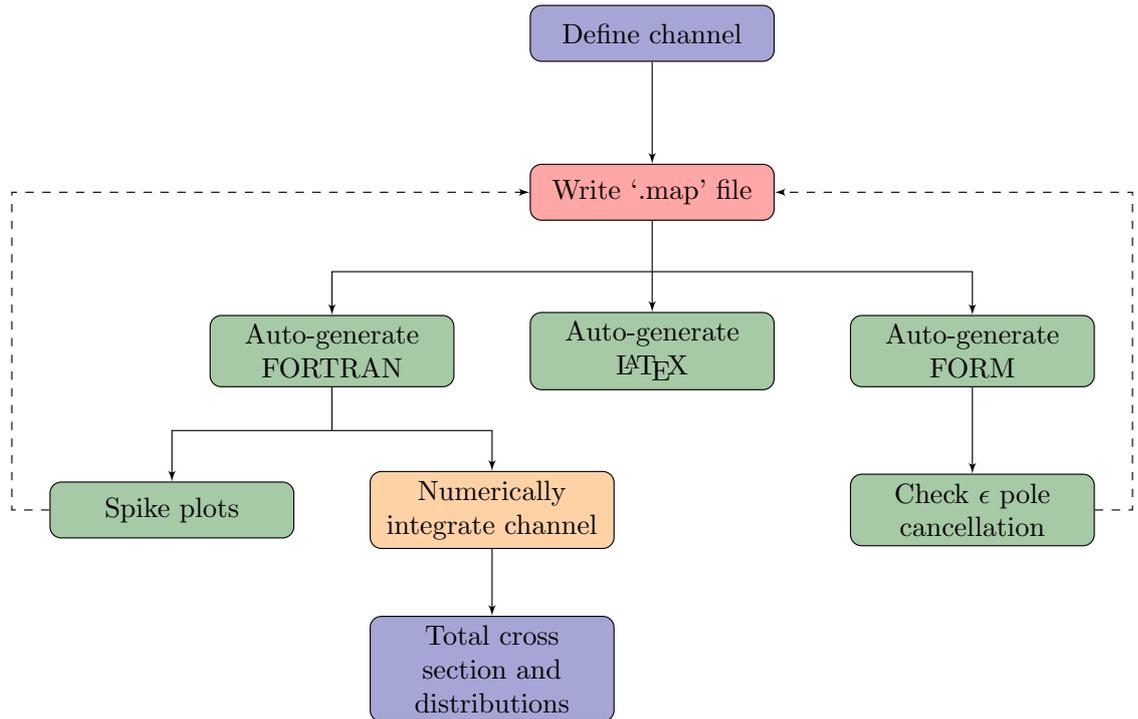


Figure 4.7: A flowchart of the processing required to obtain numerical results using NNLOJET. The red components represent a time consuming process in terms of the implementation. The orange components indicate the numerically time consuming contributions. The dashed arrows represent a feedback loop where the construction of the ‘.map’ file is determined from the quality of the spike plots and/or the cancellation of ϵ poles.

space and/or there are no explicit ϵ poles remaining, the subtraction term can then be added to the auto-generated integration routines. This can then be numerically integrated, usually in conjunction with many other channels, to generate a total cross section and differential distributions for a set of physical cuts.

4.6.3 Channel breakdown of Z +jet at NNLO

All of the NNLO corrections to Z +jet were computed using the formalism outlined in the previous section. Clearly given the vast number of colour configurations and initial states it would not be appropriate to review all of the contributions in detail. Nevertheless we shall discuss the notation used and the various contributions that are required in the full NNLO computation.

The matrix element contributions to Z +jet can be classified into 3 categories

- B-type - matrix elements that contain a single $q\bar{q}$ pair,

- C-type - matrix elements that contain 2 $q\bar{q}$ pairs that do not necessarily have to be the same flavour,
- D-type - matrix elements that contain 2 $q\bar{q}$ pairs of identical flavour quarks.

In principle there are also pure gluon matrix elements contributing to Z +jet production which first appear at NLO [152]. The Z boson couples to a closed fermion loop in these contributions. However these are a negligible contribution to the full cross section and therefore can be safely ignored.

Within each category of matrix element, we can further classify the matrix element by the number of loops in the contribution and the number of gluons. For example, a leading colour B-type matrix element with M gluons and L loops would be denoted by,

BMgLZ.

Clearly there are many colour contributions with the same parton content. To further distinguish the colour contributions, we introduce additional labels t , which denote contributions which are subleading in N , and h , which denote contributions which are proportional to n_f . A contribution from C-type matrix element subleading in N with M gluons and L loops would be expressed as

CtMgLZ.

For sufficiently complicated colour structures it is possible to have sub-subleading colour contributions in both N , n_f or a mixture of the two. This is denoted by multiple instances of t , h , or a combination of these.

Using this notation, the full Z +jet calculation in the NNLOJET framework is generated from the following Maple input file,

```
LO:= [
  [B1g0Z, [qb, g, q, lp, lm], 1]
]:

R1:= [
  [B2g0Z, [qb, g, g, q, lp, lm], 1],
  [Bt2g0Z, [qb, gt, gt, q, lp, lm], -1/nc**2],
```

```
[ C0g0Z , [ qb , Q, Qb, q , lp , lm ] , 1 / nc ] ,
[ D0g0Z , [ qb , q , qb , q , lp , lm ] , - 1 / nc ** 2 ]
]:
```

```
V1:= [
[ B1g1Z , [ qb , g , q , lp , lm ] , 1 ] ,
[ Bt1g1Z , [ qb , g , q , lp , lm ] , - 1 / nc ** 2 ] ,
[ Bh1g1Z , [ qb , g , q , lp , lm ] , nf / nc ]
]:
```

```
RR:= [
[ sB3g0Z , [ qb , g , g , g , q , lp , lm ] , 1 ] ,
[ sBt3g0Z , [ qb , g , g , g , q , lp , lm ] , - 1 / nc ** 2 ] ,
[ sBtt3g0Z , [ qb , gt , gt , gt , q , lp , lm ] , ( nc ** 2 + 1 ) / nc ** 4 ] ,
[ sC1g0Z , [ qb , g , Q, Qb, q , lp , lm ] , 1 / nc ] ,
[ sCt1g0Z , [ qb , g , Q, Qb, q , lp , lm ] , 1 / nc ** 3 ] ,
[ D1g0Zc , [ qb , g , q , qb , q , lp , lm ] , - 1 / nc ** 2 ] ,
[ Dt1g0Z , [ qb , g , q , qb , q , lp , lm ] , + 1 / nc ** 4 ]
]:
```

```
RV:= [
[ sB2g1Z , [ qb , g , g , q , lp , lm ] , 1 ] ,
[ sBt2g1Z , [ qb , g , g , q , lp , lm ] , - 1 / nc ** 2 ] ,
[ sBtt2g1Zc , [ qb , gt , gt , q , lp , lm ] , - 1 / nc ** 2 ] ,
[ sBttt2g1Z , [ qb , gt , gt , q , lp , lm ] , + 1 / nc ** 4 ] ,
[ sBh2g1Z , [ qb , g , g , q , lp , lm ] , nf / nc ] ,
[ sBth2g1Z , [ qb , gt , gt , q , lp , lm ] , - nf / nc ** 3 ] ,
[ sC0g1Z , [ qb , Q, Qb, q , lp , lm ] , 1 / nc ] ,
[ Ct0g1Z , [ qb , Q, Qb, q , lp , lm ] , 1 / nc ** 3 ] ,
[ Ch0g1Z , [ qb , Q, Qb, q , lp , lm ] , nf / nc ** 2 ] ,
[ D0g1Z , [ qb , q , qb , q , lp , lm ] , + 1 / nc ** 2 ] ,
[ Dt0g1Z , [ qb , q , qb , q , lp , lm ] , + 1 / nc ** 4 ]
]:
```

```

VV:= [
  [B1g2Z , [ qb , g , q , lp , lm ] , 1 ] ,
  [Bt1g2Z , [ qb , g , q , lp , lm ] , -1/nc**2 ] ,
  [Btt1g2Z , [ qb , g , q , lp , lm ] , 1/nc**4 ] ,
  [Bh1g2Z , [ qb , g , q , lp , lm ] , nf/nc ] ,
  [Bht1g2Z , [ qb , g , q , lp , lm ] , nf/nc**3 ] ,
  [Bhh1g2Z , [ qb , g , q , lp , lm ] , nf**2/nc**2 ] ,
  [fake_C0g2Z , [ qb , Q , Qb , q , lp , lm ] , 1/nc ] ,
  [fake_Ct0g2Z , [ qb , Q , Qb , q , lp , lm ] , 1/nc**3 ] ,
  [fake_D0g2Z , [ qb , q , qb , q , lp , lm ] , 1/nc**1 ] ,
  [fake_Dt0g2Z , [ qb , q , qb , q , lp , lm ] , 1/nc**3 ]
]:

```

where *LO*, *R1*, *V1*, *RR*, *RV*, *VV* are lists of terms included in their respective calculations. There are some additional points to note from this list. Firstly, some of the matrix elements begin with an ‘s’. This is to denote that in some initial state configurations the final state is symmetrised. For example, *sB3g0Z* would imply that in the configuration where two of the gluons are in the initial state, the individual colour contributions are given by,

$$sB_3^0(3_q, \hat{1}_g, \hat{2}_g, 4_g, 5_{\bar{q}}) = \frac{1}{2} \left(B_3^0(3_q, \hat{1}_g, \hat{2}_g, 4_g, 5_{\bar{q}}) + B_3^0(5_{\bar{q}}, 4_g, \hat{2}_g, \hat{1}_g, 3_q) \right). \quad (4.6.1)$$

This is so that for *gg* initiated processes we recover a full line reversal symmetry. This is convenient for the construction of subtraction terms, despite the small cost of numerically having to evaluate the matrix element twice. If one of the initial state partons is not a gluon however, these matrix element return their unsymmetrised value. For example,

$$sB_3^0(\hat{1}_q, \hat{2}_g, 3_g, 4_g, 5_{\bar{q}}) = B_3^0(\hat{1}_q, \hat{2}_g, 3_g, 4_g, 5_{\bar{q}}). \quad (4.6.2)$$

Some C-type matrix elements also are denoted by ‘s’. This is similar to the B-type case, however instead we exchange the momentum of a final state $Q\bar{Q}$ pair. For example,

$$sC_1^0(\hat{1}_q, \hat{2}_g, 3_{\bar{Q}}, 4_Q, 5_{\bar{q}}) = \frac{1}{2} \left(C_1^0(\hat{1}_q, \hat{2}_g, 3_{\bar{Q}}, 4_Q, 5_{\bar{q}}) + C_1^0(\hat{1}_q, \hat{2}_g, 4_{\bar{Q}}, 3_Q, 5_{\bar{q}}) \right). \quad (4.6.3)$$

This is again a convenient choice for constructing subtraction terms. If there is not

an identical $Q\bar{Q}$ pair in the final state then no exchange of momenta is applied. For example,

$$sC_1^0(\hat{1}_q, 3_g, \hat{2}_{\bar{Q}}, 4_Q, 5_{\bar{q}}) = C_1^0(\hat{1}_q, 3_g, \hat{2}_Q, 4_Q, 5_{\bar{q}}). \quad (4.6.4)$$

In some cases a subtraction term must be integrated which does not have any corresponding matrix element. For example, the gg initiated real corrections to Z +jet have no corresponding virtual matrix element. In these cases the antenna functions introduced cancel against mass factorisation, as outline in chapter 2. In most cases the driver routines generate a corresponding subtraction term as a separate channel where no corresponding matrix element is called. For example, the $gg \rightarrow q\bar{q}$ leading colour real subtraction term is integrated out in the channel,

```

if (ip(85).and.coloursel(1)) then
  iproc = 85
  nfB1 = 1
  ip1 = 0
  ip2 = 0
  wt = ggBlg1ZTNLO(ix , xri1 , xri2 , 1 , 2 , 3 , 4 , 5 , 0 d0 , 0 d0)
  wtsum = wtsum + wt
endif

```

where the overall factors and PDFs match the corresponding real corrections. In some cases however this is not always possible to do uniquely. This is as a result of the problem outlined in section 2.6, where in some two quark pair processes the flavour of the partons decouple from the flavour of the PDFs. Hence there are some ‘fake’ routines to deal with this specific cases that are unique to calculations involving vector bosons.

4.7 Numerical optimisation

As demonstrated in Tab. 4.1, a single calculation Z +jet up to NNLO has a significant CPU overhead. It is important that the calculation is optimised to be as efficient as possible. This section will discuss some of the optimisations implemented within the code, primarily focusing on double real contributions.

4.7.1 Storage element

For real and double real phase spaces, there are multiple possible combinations of factorised phase spaces that must be evaluated. For large final state multiplicities the number of possible reduced phase spaces that must be evaluated can be large. To demonstrate the scaling of this problem, consider a process $pp \rightarrow n$ partons at leading order. For the real contribution, this implies there are $n + 1$ partons in the final state. The number of single unresolved mapped phase spaces will be denoted by the function $S(n)$, where we must separately consider contributions from the initial-initial, initial-final and final-final mapping functions.

For the possible initial-initial mappings, there are $n + 1$ possible phase spaces that can be evaluated. For initial-final mappings, there are $n(n + 1)$ possible combinations. This is because there are $n + 1$ possible choices for the unresolved parton, followed by n choices for the final state hard radiator. There are two possibilities for the initial state hard radiator, increasing the combinations by a factor of 2, however the initial-final mapping function is symmetric under the exchange of the final state partons, hence this reduces the number of possibilities by a factor of 2 again. Using a similar counting argument for the final-final mappings, there are,

$$\frac{(n - 1)n(n + 1)}{2}, \quad (4.7.1)$$

possible mappings. The factor of $1/2$ is included because final-final mapping is symmetric under a line reversal of the entire mapping.

Adding together all of the number of possible combinations gives the function,

$$S(n) = (n + 1) \left[1 + \frac{n(n + 1)}{2} \right]. \quad (4.7.2)$$

For the double real phase space one has to include the number of possible combinations of phase spaces arriving from colour connected double unresolved mappings. This will be denoted by $D(n)$, where there are contributions from double unresolved initial-initial, initial-final and final-final mappings.

For the initial-initial double unresolved mapping function there are,

$$\frac{(n + 1)(n + 2)}{2}, \quad (4.7.3)$$

possible phase spaces. Again the factor of $1/2$ is included because the double unresolved initial-initial mapping function is symmetric under the exchange of the unresolved partons. In the case of double unresolved initial-final mappings, there are,

$$\frac{2n(n+1)(n+2)}{3!}, \quad (4.7.4)$$

possible phase spaces. This is because the double unresolved initial-final mapping function is symmetric under the permutation of the final state partons, hence there are $3!$ identical combinations of phase spaces compared to the naïve counting. The additional factor of 2 must be included because there are two possible initial state hard radiators.

For the double unresolved final-final case, there are,

$$\frac{(n-1)n(n+1)(n+2)}{2}, \quad (4.7.5)$$

possible phase spaces. The additional factor of 2 is required because the double unresolved final-final mapping is symmetric under a line reversal of the entire mapping.

Adding together all of these possible combinations gives,

$$D(n) = \frac{(n+1)(n+2)}{2} \left[\frac{3+2n}{3} + n(n-1) \right]. \quad (4.7.6)$$

For a real phase space contribution to a given process, the number of possible combinations of reduced phase spaces $R(n)$ is simply the number of possible single unresolved phase spaces,

$$R(n) = S(n). \quad (4.7.7)$$

As discussed in section 2.4.1, there are 3 possible types of factorised phase spaces for a double real contribution,

- single unresolved mappings onto an $(n+1)$ -parton phase space,
- double unresolved mappings onto an n -parton phase space,
- two iterated single unresolved mapping functions which map onto an n -parton phase space.

This implies that the number of possible combinations of reduced phase spaces for the

| n | $R(n)$ | $RR(n)$ |
|-----|--------|---------|
| 0 | 1 | 9 |
| 1 | 4 | 65 |
| 2 | 12 | 390 |
| 3 | 28 | 1685 |
| 4 | 55 | 5611 |

Table 4.2: The number of combinations of reduced phase spaces for real corrections ($R(n)$) and double real corrections ($RR(n)$) to the process $pp \rightarrow n$ partons, where there are n partons in the leading order contribution.

double real contribution, $RR(n)$, is given by,

$$RR(n) = S(n+1)[S(n)+1] + D(n). \quad (4.7.8)$$

The number of possible combinations of reduced phase spaces up to $n = 4$ are given in Tab. 4.2. From Tab. 4.2 we can see that for Z +jet production, where $n = 1$, there are 65 distinct reduced phase spaces that must be evaluated for each double real event. For a full NNLO calculation it is likely that all of these phase spaces will be computed. Since the code is factorised on the basis of initial state and colour contribution, many of these phase spaces will be required multiple times. This is a huge bottleneck on the double real calculation and it is more appropriate to store as much of the information as possible.

The phase space storage element consists of a large array, where each element of the array is subsequently a pointer to *all* of the kinematic information for a given factorised phase space. This includes partonic information, such as invariants and spinors, and also final state observables.

The first time the code comes across a particular mapping, it will generate all of the kinematic information and then store them in the array. A flag will be set to indicate that the information has been stored for that particular mapping. For each subsequent time that the code needs to use this particular mapping it will recall the information from the array, saving a huge amount of resources that scales extremely well for a large number of processes. Once the evaluation of the full phase space is complete, the array is flushed by simply setting all of the flags to false. The next event in the series will

then overwrite the pointers in the array.

For calculations involving a dynamic mass factorisation scale choice, the sets of PDFs can also be stored in a similar fashion to above. Dynamic scale choices need to be reevaluated from the final state observables for each contribution to the subtraction term, hence each term would require a new evaluation of the PDFs for that particular scale. These can be stored to reduce the number of evaluations of the PDFs.

4.7.2 Colour sampling

As discussed in section 4.6.2, the NNLOJET channels are divided both into initial state and colour configurations. There are a large number of subleading colour channels at NNLO which are a numerically small contribution to the calculation. To avoid costly evaluations of subleading colour channels as much as possible, a multichannel calculation was implemented that samples the larger colour channels favourably over the subleading colour channels. This is performed by transforming the integral such that,

$$\int_{d\sigma} \sum_{\mathcal{C}} d\sigma_{\mathcal{C}} = \int_{d\sigma} \int_0^1 dx_{\mathcal{C}} \sum_i \Theta(x_{\mathcal{C},i+1} - x_{\mathcal{C}}) \Theta(x_{\mathcal{C}} - x_{\mathcal{C},i}) \frac{d\sigma_{\mathcal{C}}}{x_{\mathcal{C},i+1} - x_{\mathcal{C},i}}, \quad (4.7.9)$$

where $d\sigma_{\mathcal{C}}$ denotes a specific colour channel and Θ denotes the Heaviside function. $x_{\mathcal{C}}$ denotes a new random number introduced into the calculation for the Monte Carlo integrator to sample. The numbers $x_{\mathcal{C},i}$ and $x_{\mathcal{C},i+1}$ denote the boundaries over which the random number $x_{\mathcal{C}}$ are sampled. To sample a particular channel more often, $x_{\mathcal{C},i+1} - x_{\mathcal{C},i}$ should be made larger.

Unlike with the other random numbers, it is *not* advantageous to use the VEGAS algorithm to perform importance sampling on $x_{\mathcal{C}}$. The VEGAS algorithm adapts large changes in the derivative of the integrand [149]. The $x_{\mathcal{C}}$ random number is effectively made discrete through the use of the Heaviside functions. The VEGAS algorithm would favourably sample around the boundaries between colour channels, which is not what is desired.

To perform the multichannel integration over colour channels, the colour channels $d\sigma_{\mathcal{C}}$ must be defined. Each colour level for a specific initial state configuration is ordered according to the highest powers of N and n_f in their overall colour factor $\mathcal{C}(N, n_f)$.

| Channel id | Initial states | | | |
|------------|-------------------|-------------------|------------------------|----------------------------------|
| | $qg + \bar{q}g$ | $q\bar{q}$ | gg | $qQ + q\bar{Q} + \bar{q}\bar{Q}$ |
| 1 | 1 | 1 | n_f | $\frac{1}{N}$ |
| 2 | $\frac{n_f}{N}$ | $\frac{n_f}{N}$ | $\frac{N^2+1}{N^4}n_f$ | $\frac{1}{N^3}$ |
| 3 | $\frac{1}{N^2}$ | $\frac{1}{N^2}$ | N/A | N/A |
| 4 | $\frac{n_f}{N^3}$ | $\frac{n_f}{N^3}$ | N/A | N/A |
| 5 | $\frac{1}{N^4}$ | $\frac{1}{N^4}$ | N/A | N/A |

Table 4.3: A table of colour factors which are sampled for each initial state configuration separated by colour channel for the Z +jet double real integration. The importance of the sampling is in reverse numerical order (i.e. channel 1 is sampled most frequently and channel 5 is sampled the least often). In channels 3, 4 and 5 there are no longer any colour contributions for the gg and the $qQ + q\bar{Q} + \bar{q}\bar{Q}$ initial states, hence only the qg and gq channels are sampled for these channels. Contributions proportional to n_f are divided into n_{up} and n_{down} processes for Z +jet production.

This is dictated by the parameter p ,

$$p = \text{pow}(N) + \frac{1}{2}\text{pow}(n_f), \quad (4.7.10)$$

where ‘pow’ denotes the highest power of their respective factors in $\mathcal{C}(N, n_f)$. The factor of $1/2$ is a choice to indicate that typically n_f contributions are smaller contributions to the final result in comparison to equivalent powers of N .

Once the ordering of the colour levels has been established for every initial state configuration, the colour channels are taken to be a sum over all of the initial states for each position in the ordering of p . This means that in each colour channel, different initial state contributions can have a different overall colour factor. For example, in the first colour channel for the double real Z +jet production at NNLO the $q\bar{q}$ initial state has an overall colour factor of 1, whereas the leading qQ contribution has an overall colour factor of $1/N$. Essentially the hierarchy of importance sampling is dictated by a particular colour contribution *relative* to the rest of the colour contributions with that particular initial state configuration. The colour channels selected for the double real Z +jet integration are listed in Tab. 4.3.

In principle we could have constructed a multichannel integration on the basis of there being a one-to-one correspondence between the initial state and colour channels

and those used in the multichannel integration. This could potentially bias the results obtained however. A particular initial state configuration may only be a small contribution to the overall cross section, for example, but differentially it could make up a large contribution to the result for a particular observable. For example, the qQ initial state is an extremely small contribution to the calculation at small transverse momentum, where the vast majority of the weight of the cross section resides. However at large transverse momentum, where the cross section is small, it is the dominant initial state contribution. This would potentially result in an undersampling of the distribution at large transverse momentum.

For the Z +jet calculation, the boundaries are fixed such that each subleading colour contribution is sampled three times less often than the previous colour channel. Clearly this is somewhat ad-hoc and work is on-going to improve this such that we adapt the boundaries depending on the calculation in question. Nevertheless, even with this naïve boundary choice the numerical performance gain compared to no colour sampling is approximately a factor of four.

4.7.3 Reweighting

As discussed in the previous section, for observables with a sharply falling weight (e.g. transverse momentum distributions) the tail of the distribution can be undersampled. One solution to this problem is to pass back to VEGAS a different weight, namely,

$$p_T^n \cdot \sigma, \quad (4.7.11)$$

where n is a positive number chosen such that the distribution,

$$p_T^n \cdot \frac{d\sigma}{dp_T}, \quad (4.7.12)$$

is approximately uniform as a function of p_T . This means that the VEGAS algorithm will attempt to sample the integrand across a wide range of p_T .

The reweighting procedure should be used with extreme caution. Whilst it can improve one particular distribution, it can have a severe negative impact on other distributions that do not have the same sharply falling weight. For example, rapidity distributions have a very different shape compared to transverse momentum distribu-

tions. Reweighting to improve a transverse momentum distribution would lead to an undersampling of the central region of the rapidity distribution, leading to a much slower convergence. Typically reweighting should be reserved for integrals dedicated to one particular observable.

4.7.4 Multiscale integration

As outlined in section 4.3, the scale dependence of the calculation is predictable and can be inferred from another scale choice. Numerically the shape and distribution of the weight of the integrand is largely the same between different (reasonable) scale choices. This means that it is perfectly reasonable to use VEGAS to sample a particular scale choice and use the same set of random numbers to sample other scale choices simultaneously.

In practice this is implemented by promoting scale dependent terms in the calculation (e.g. PDFs, $\alpha_s(\mu_R^2)$, etc) to have an additional dimension. Each element of the new dimension corresponds to a different choice of scales for μ_F and μ_R . These can then be combined with the scaleless terms (e.g. tree level matrix elements, unintegrated antennae, etc) to generate the full result for an arbitrary number of scales simultaneously. This effectively leads to a huge performance boost given that the numerically expensive contributions to the calculation are from the scaleless terms. Indeed, using the NNLOJET code to calculate 7 scale choices simultaneously takes approximately the same amount of CPU time to an identical calculation involving just 1 scale. Each scale does not need to be proportional to a common observable either, one could run fixed scale choices with dynamical scale choices simultaneously.

This formalism could be extended to support multiple different PDF sets too, where different PDF sets would correspond to different ‘scales’. With many PDF sets running simultaneously there could be a negative impact on the performance of the code however. If this proves to be the case, it would be advantageous to use `applgrid` [153], `fastNLO` [154] or `n-tuples` [155]. These codes can be used to change the PDF sets *a posteriori*. All of these tools were designed for NLO calculations however and would require further efforts to use them for a non-trivial NNLO computation.

4.7.5 Observable matching

A common source of statistical fluctuations in differential distributions is where the weight of the matrix element and the subtraction term in an unresolved configuration lie across a bin edge. The discussion of chapter 5 will focus on a treatment of these fluctuations following integration. We can however reduce the size of the fluctuations by matching the observables of the subtraction term to those derived from the full phase space below a certain threshold. This means that in very unresolved configurations the matrix element and subtraction term are guaranteed to lie within the same bin.

To compute the threshold required, there must be a definition for how unresolved a particular phase space configuration is. This can be inferred by analysing the invariants of the partons involved in a particular mapping function. Consider a mapping $(i, j, k) \rightarrow (I, K)$, where j is an unresolved parton. The threshold condition is defined such that,

$$\min \left(\frac{s_{ij}}{s_{ik}}, \frac{s_{jk}}{s_{ik}} \right) < \epsilon_0, \quad (4.7.13)$$

where ϵ_0 is a small dimensionless parameter. If this condition is satisfied then the observables for the subtraction term are taken from the full phase space instead of being calculated from its factorised phase space. This condition will correctly isolate both of the single collinear limits $s_{ij} \rightarrow 0$, $s_{jk} \rightarrow 0$ and also the j soft limit.

This procedure can be extended for double unresolved mapping functions. Consider the double unresolved mapping $(i, j, k, l) \rightarrow (I, L)$ where both j and k are unresolved. The threshold condition for this case is,

$$\min \left(\frac{s_{ijk}}{s_{il}}, \frac{s_{jkl}}{s_{il}}, \frac{s_{ij} + s_{kl}}{s_{il}}, \frac{s_{ik} + s_{jl}}{s_{il}}, \frac{s_{ij} + s_{ik} + s_{kl}}{s_{il}}, \right. \\ \left. \frac{s_{ik} + s_{ij} + s_{jl}}{s_{il}}, \frac{s_{kl} + s_{ij} + s_{jl}}{s_{il}}, \frac{s_{jl} + s_{ik} + s_{kl}}{s_{il}} \right) < \epsilon_0, \quad (4.7.14)$$

which correctly captures all of the possible triple collinear, double collinear and soft collinear limits in a colour connected double unresolved mapping. The double soft limit will automatically be accommodated by one of the above constraints.

It should be noted Eq. (4.7.14) cannot correctly match configurations in single unresolved limits. In general this would be extremely challenging to implement because it would require matching to the many possible $(n + 1)$ -parton mapped phase spaces.

Nevertheless, the double unresolved limits are generally a far greater contribution to the statistical fluctuations than single unresolved limits.

As with the full and reduced technical cuts, it is important that the parameter ϵ_0 is not set too high. If it is set inappropriately, the distributions will be artificially smeared and will not represent the physical result. A typical choice for ϵ_0 is $3 \cdot y_0$, meaning that only limits that are extremely close to the technical limit are matched in this way.

Chapter 5

Analysing data from NNLO calculations

Following the running of the NNLOJET code, the output of many differential distributions and total cross sections need to be combined to extract the physical result. There is a large degree of flexibility in how to do the combination which, in general, has to balance the need to regulate bin-to-bin fluctuations and at the same time remain in a regime which is still statistically valid. This chapter will discuss one approach for combining large quantities of data from NNLO calculations such that the physical results can be extracted in a meaningful fashion.

5.1 Statistical Outliers

When using the antenna subtraction method with a phase space with real emissions, the observables of the full matrix element and individual contributions to a subtraction term are evaluated in different phase spaces. The observables of the matrix element are evaluated using the full partonic phase space, whereas the observables of the subtraction term are evaluated using the factorised phase space appropriate to the mapping in question. In the absolute unresolved limits, the observables of the matrix element and subtraction term must match. However, if the phase space is a small, finite distance away from the singularity then there can be a small difference between the sets of observables. This can produce two forms of outliers,

- the matrix element (subtraction term) is evaluated within the region defined by the fiducial cuts, while the subtraction term (matrix element) is evaluated outside the region. This generates an artificially large weight close to the kinematical cut boundary which is not representative of the true integrand.
- In differential results, the matrix element and subtraction term lie either side of a bin boundary, leading to bin-to-bin fluctuations.

At NLO, the solution to this problem is straightforward. The technical cut should be tuned to be as high as possible, preventing the integration routines from sampling too close to the phase space singularities. Clearly this needs to be balanced by the fact that there must be no residual technical cut dependence on the calculation. Beyond this it is simply a case of performing a sufficient number of evaluations of the integrand that the outliers are averaged away and as such become an irrelevant contribution to the final result.

At NNLO the solution is not as simple. The double real subtraction term consists of many contributions, with an internal cancellation happening in every divergent configuration of the phase space. As discussed in section 4.7, for $pp \rightarrow 3$ partons there are 65 possible combinations of factorised phase spaces that can be evaluated for a single double real event. For a calculation involving all initial states and colour combinations, it is likely that all of these phase spaces will be generated. The rate of outlying events is therefore significantly higher than for NLO calculations. It is not feasible to simply run for a large enough number of events to make outliers an irrelevant contribution to the calculation, a more sophisticated solution is required.

5.2 Weighted averaging

A common approach to deal with outliers is to use a weighted average, both in the total cross section and in the differential distributions. Consider a double real integrand evaluated $\mathcal{O}(10^9)$ times divided into N runs, each containing M events. For a single run (of M events) denoted by i , the unweighted average of the function $f(\vec{x})$ is given by,

$$\sigma_i(M) = \sum_{j=1}^M \frac{f(\vec{x}_j)}{M}, \quad (5.2.1)$$

with corresponding unweighted statistical uncertainty $\delta\sigma_i(M)$. The weighted average of N such runs is given by,

$$\sigma(N, M) = \frac{\sum_{i=1}^N \omega_i(M) \sigma_i(M)}{\sum_{i=1}^N \omega_i(M)}, \quad (5.2.2)$$

where $\omega_i(M)$ is the weight of a particular run i , given by,

$$\omega_i = \frac{1}{\delta\sigma_i^2(M)}. \quad (5.2.3)$$

A run containing outliers is typically accompanied with a larger error and, thus, enters the weighted average with a smaller weight ω_i . In particular, given a fixed number of integrand evaluations ($N \times M$ fixed):

- For large M , small N , the $\sigma_i(M)$ have a Gaussian distribution but the impact of outliers is not taken care adequately, leading to a large uncertainty $\delta\sigma(N, M)$.
- For small M , large N , the $\sigma_i(M)$ do not have a Gaussian distribution and are dominated by outliers but the statistical weighting reduces the effect of outliers leading to a small uncertainty $\delta\sigma(N, M)$. In this case, a wrong value for $\sigma(N, M)$ is determined but with a small $\delta\sigma(N, M)$.

To optimise the efficiency of the numerical integration, a compromise between N and M must be found while guaranteeing the correctness of the final answer. For calculating the double real contributions to the Z +jet calculation at NNLO N is fixed to be $\mathcal{O}(10^4)$ and M is fixed to be $\mathcal{O}(10^5)$. The cross section $\sigma(N/k, k \cdot M)$ can then be defined as a function of k , as shown in Fig. 5.1.

As expected, as k increases, the statistical error grows. On the other hand, as k decreases, there is a sizeable shift in the central value which amounts to a difference of almost 1% at $k = 1$ with respect to the value of the cross section in the stable region.

As a consistency check, Fig. 5.2 shows the total cross section computed from the integral of the Z boson rapidity distribution as a function of k . As in Fig. 5.1, there is a sizeable k dependence. In particular, for $k < 100$ the statistical errors are small, but the central value is off by a large amount. For $k > 100$, the central value is stable, but the statistical errors increase.

The impact of performing a weighted average can be assessed on the Z boson rapidity distribution itself. Fig. 5.3 shows that as k increases, the bin-by-bin statistical

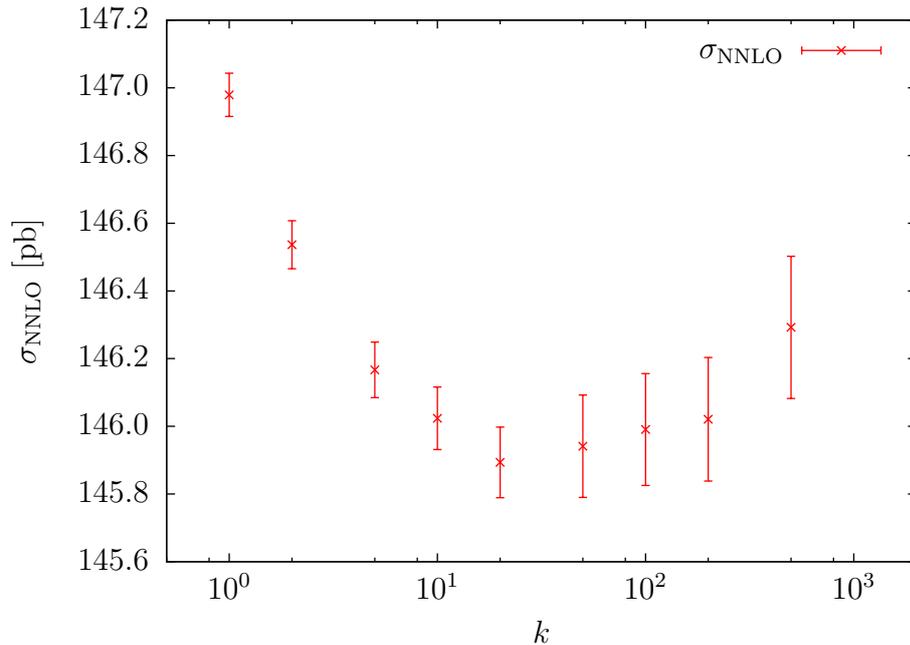


Figure 5.1: The total cross section $\sigma(N/k, k \cdot M)$ at NNLO as a function of k , the number of runs combined to form a pseudorun with $k \cdot M$ events. The cross section is then the weighted average of the pseudoruns. The errors on the individual data points are statistical.

errors also increase (as expected). For $k < 5$, the central values for different values of k disagree by a sizeable amount, resulting in an incorrect value for the total cross section evaluated from the integral of this distribution, as shown in Fig. 5.2.

Similar conclusions can be drawn from transverse momentum distribution of the Z boson shown in Fig. 5.4. Above the p_T^{jet} cut, the p_T^Z distribution is stable as a function of k , even at low values. This demonstrates that there are fewer outliers in this distribution compared to the rapidity distribution of the Z boson, as expected, and the statistical errors are more reliable at small k .

Below the p_T^{jet} cut there is a sizeable shift in the central value for small k . This can be understood because this region probes configurations where the extra partonic radiation present at NLO and NNLO can compensate the transverse momentum of the leading jet, generating a Z boson transverse momentum below the p_T^{jet} cut—the well known Sudakov shoulder phenomenon [156]. This unusual configuration is very sensitive to the jet definition and is prone to generating outliers because of the close proximity to the physical cuts within these phase space configurations.

The conclusion to be drawn from these results suggests that a safe value for k

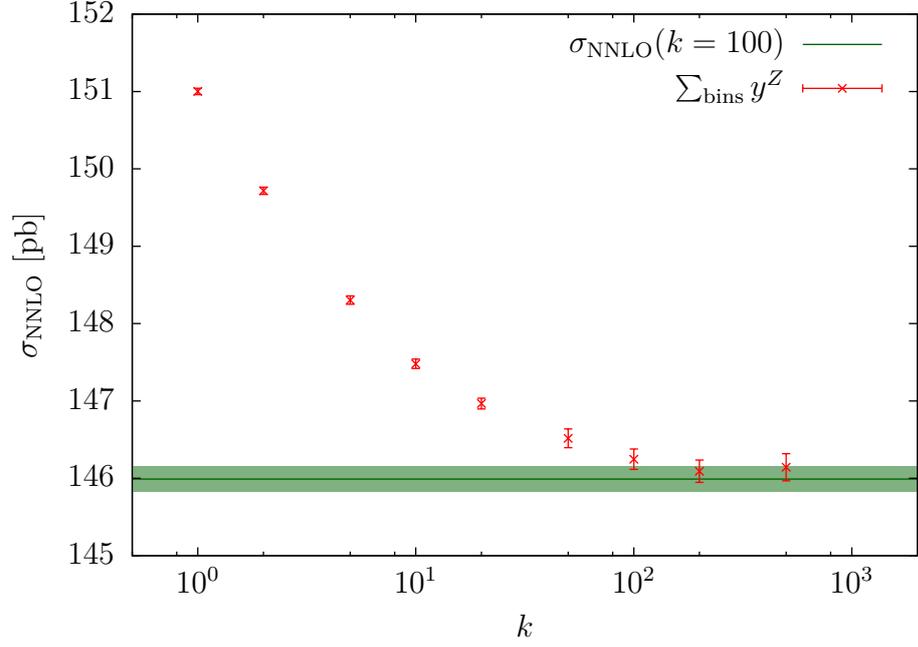


Figure 5.2: The integral of the y^Z distribution, $\sigma(N/k, k \cdot M)$ at NNLO as a function of k . The errors on the individual data points are statistical. The green band denotes the total cross section evaluated at $k = 100$ with its statistical uncertainty.

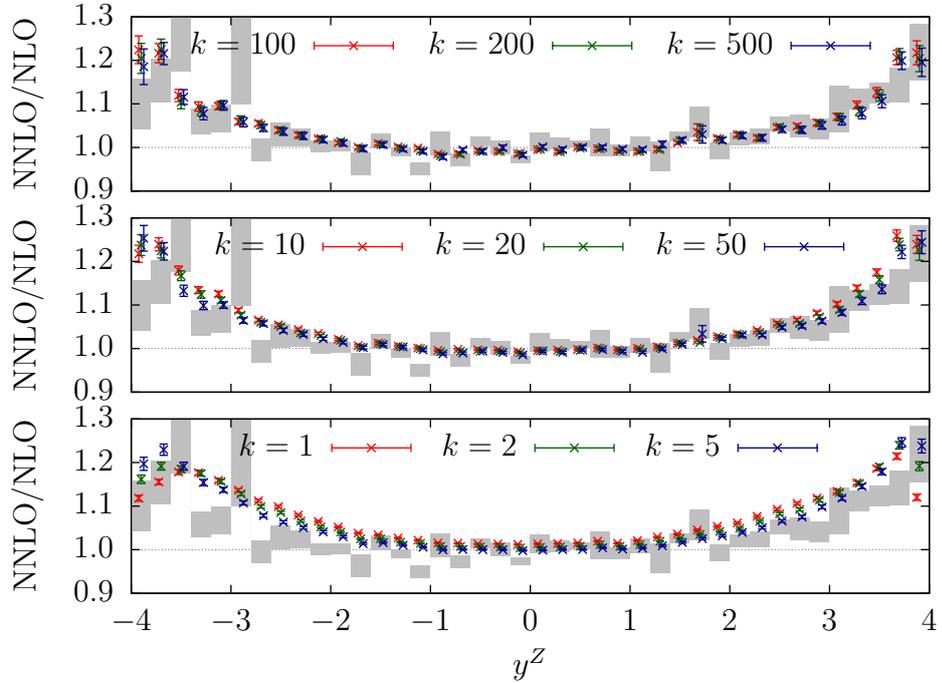


Figure 5.3: The ratio of the NNLO against the NLO prediction for individual bins of the rapidity distribution of the Z boson for different values of k . The grey bands are the unweighted result. The errors on the individual data points are statistical.

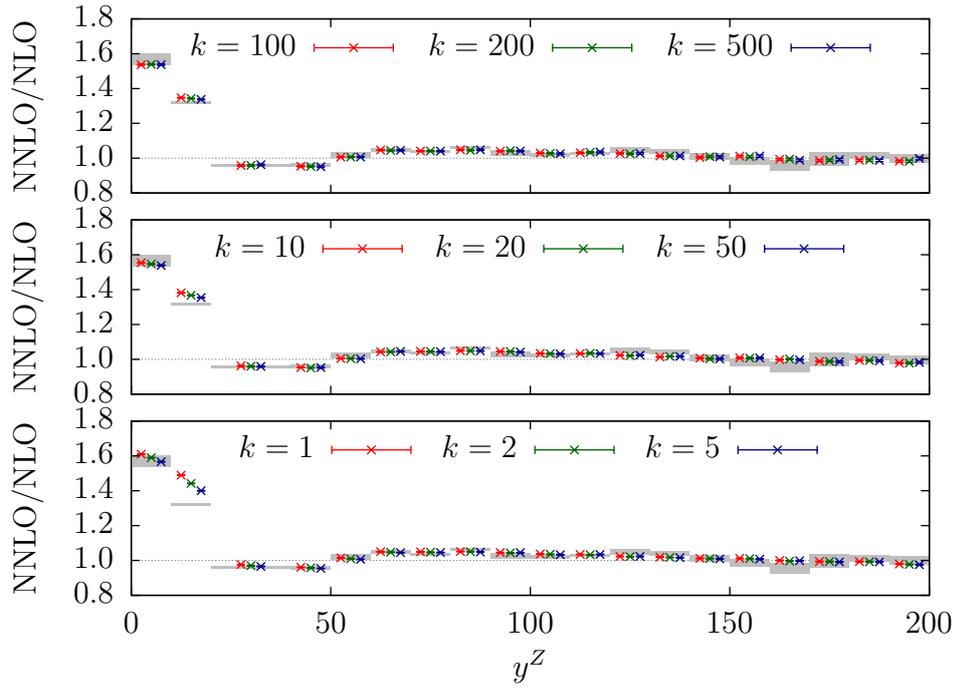


Figure 5.4: The ratio of the NNLO against the NLO prediction for individual bins of the p_T distribution of the Z boson for different values of k . The grey bands are the unweighted result. The errors on the individual data points are statistical.

is $\mathcal{O}(100)$, corresponding to $\mathcal{O}(10^7)$ events per pseudorun. For these k values the total cross section (computed directly and from integrating the distributions) and the distributions themselves are insensitive to the precise choice of k . The final results will not be shifted by outliers and the the statistical uncertainties are under control. Therefore $k = 100$ is used for all the differential distributions presented in chapter 6.

Chapter 6

Phenomenological results

This chapter will review new NNLO results of phenomenological interest for Z boson production in association with hard hadronic events. The discussion in section 6.1 will focus on establishing the size of the NNLO corrections to the Z +jet production using the antenna subtraction formalism using representative cuts on the final state jets and leptons.

In section 6.2 the NNLO corrections to the single and double-differential transverse momentum distributions for fully inclusive Z/γ^* production is presented, including the leptonic decay. We also compute the same distributions normalised to the cross sections for inclusive Z/γ^* production, i.e. integrated over the transverse momentum of the lepton pair. The NNLO prediction is compared to the 8 TeV data set from the ATLAS [143] and CMS [145] collaborations, which both observed a tension between data and NLO theory predictions, using the experimental cuts and binning.

6.1 NNLO corrections to Z +jet production at the LHC

For the numerical computations in this section, we use the NNPDF2.3 parton distribution functions [64] with the corresponding value of $\alpha_s(M_Z) = 0.118$ at NNLO, and $M_Z = 91.1876$ GeV. The same PDF set and value for $\alpha_s(M_Z)$ is used for the LO, NLO and NNLO predictions. The factorisation and renormalisation scales are chosen to be $\mu \equiv \mu_F = \mu_R = M_Z$, with a theoretical uncertainty estimated by varying the scale

| channel | cross section [pb] | | |
|------------------|-----------------------|-----------------------|-----------------------|
| | LO | NLO | NNLO |
| qg | $53.6^{+4.5}_{-4.3}$ | $80.2^{+3.5}_{-3.2}$ | $84.0^{+2.0}_{-3.9}$ |
| $q\bar{q}$ | $27.1^{+1.5}_{-1.5}$ | $33.1^{+0.2}_{-0.6}$ | $32.0^{+1.5}_{-2.0}$ |
| $\bar{q}g$ | $22.9^{+1.7}_{-1.7}$ | $33.1^{+0.0}_{-0.4}$ | $34.9^{+2.2}_{-3.3}$ |
| gg | N/A | $-4.0^{+3.3}_{-1.9}$ | $-7.2^{+5.1}_{-3.8}$ |
| qq | N/A | $1.8^{+2.7}_{-1.8}$ | $2.0^{+3.2}_{-2.6}$ |
| $\bar{q}\bar{q}$ | N/A | $0.1^{+0.4}_{-0.3}$ | $0.1^{+0.4}_{-0.4}$ |
| total | $103.6^{+7.7}_{-7.5}$ | $144.4^{+9.0}_{-7.2}$ | $145.8^{+0.0}_{-1.2}$ |

Table 6.1: Channel breakdown of the total cross section for LO, NLO and NNLO for the scale choice $\mu = \mu_F = \mu_R = M_Z$. The theoretical uncertainty on each channel is estimated from the envelope of the $[1/2, 1, 2] \cdot M_Z$ scale choices.

choice by a factor in the range $[1/2, 2]$.

The leptons are required have pseudorapidity, $|\eta^\ell| < 5$ and the dilepton invariant mass must be close to the Z boson mass, $80 \text{ GeV} < m_{\ell\ell} < 100 \text{ GeV}$. Jets are reconstructed using the anti- k_T algorithm [87] with $R = 0.5$ and are required to have $p_T^{\text{jet}} > 30 \text{ GeV}$ and $|y^{\text{jet}}| < 3$.

Tab. 6.1 shows the total cross section found using these cuts at LO, NLO and NNLO, including a channel breakdown based on the initial state. Altogether the inclusive NNLO corrections amount to a 1% increase on the NLO cross section. For each channel there are relatively small shifts in their contribution to the total cross section going from NLO to NNLO. However, the most significant shift is in the gg initiated channel which receives a correction of nearly 80% compared to the NLO prediction. This significant shift is not surprising as this channel first appears at NLO. The gg contribution lies mainly around $y^Z \sim 0$ and will be highly relevant for fitting the gluon PDFs using the Z +jet LHC data.

Fig. 6.1 shows the inclusive leading jet transverse energy distribution in 10 GeV bins. Due to the inclusiveness of the observable, events with two or three jets with $p_T^{\text{jet}} > 30 \text{ GeV}$ and $|y^{\text{jet}}| < 3$ are also included. The relative corrections are further exposed in the lower panel of Fig. 6.1, where the NLO and NNLO K factors are presented differentially. The band shows the effect of varying $\mu = \mu_R = \mu_F$ in the

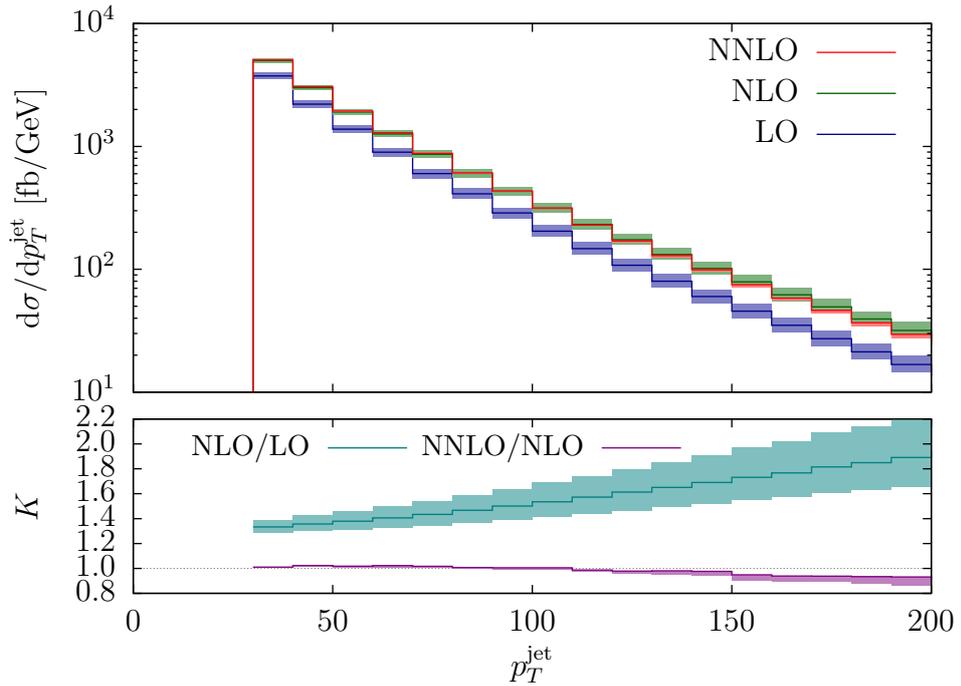


Figure 6.1: The transverse momentum distribution of the leading jet in inclusive Z +jet production in pp collisions with $\sqrt{s} = 8$ TeV at LO (blue), NLO (green), NNLO (red). The lower panel shows the perturbative K factors for NLO/LO (turquoise) and NNLO/NLO (mauve).

range $[1/2, 2] \cdot M_Z$ in the numerator while keeping $\mu = M_Z$ in the denominator. For this set of cuts and input parameters the NLO corrections increase the cross section by between 30% to 70%. At low transverse momentum the NNLO corrections are a positive correction of approximately 1%. The variation with the unphysical scales in the NNLO prediction is significantly reduced in comparison to the NLO prediction.

The rapidity distribution of the leading jet is displayed in Fig. 6.2. Note that the distribution is restricted by the requirement that $|y^{\text{jet}}| < 3$. The NLO corrections are typically 35%–40% and relatively flat. The NNLO corrections increase the cross section by approximately 1% over the whole range of y^{jet} with a significantly reduced scale dependence.

The Z boson p_T distribution in inclusive Z +jet production is shown in Fig. 6.3, where the Sudakov shoulder phenomenon [156] is observed at $p_T^Z \approx 30$ GeV. The NNLO corrections tend to stabilise the NLO result close to the Sudakov shoulder, and in fact simply represent a NLO correction to the p_T^Z distribution for Z +jet events in this region. At larger transverse momenta, the NNLO corrections increase the prediction by approximately 1%.

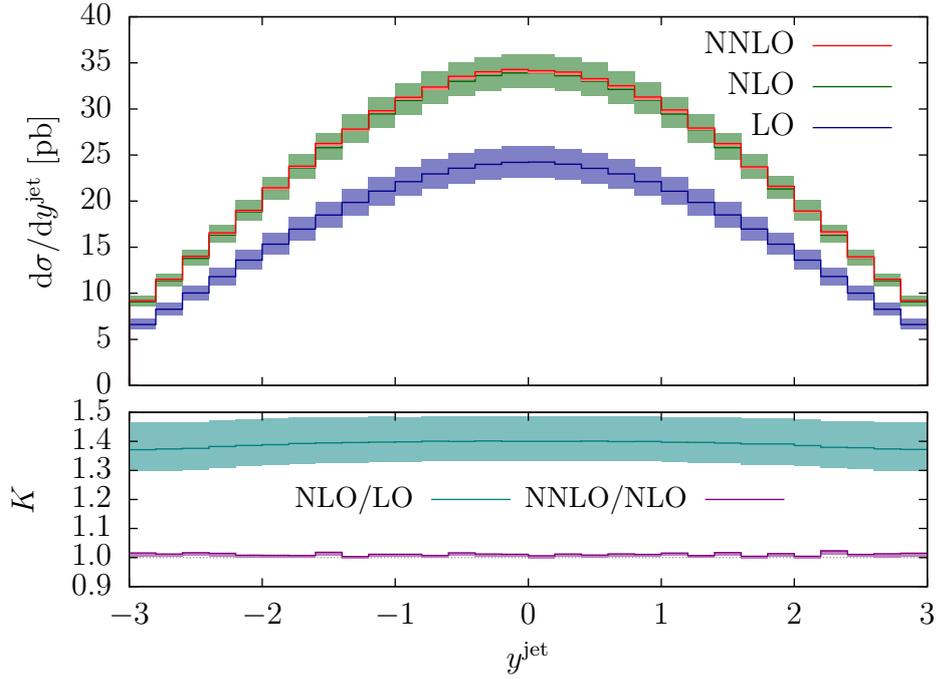


Figure 6.2: The rapidity distribution of the leading jet in inclusive Z +jet production in pp collisions with $\sqrt{s} = 8$ TeV at LO (blue), NLO (green), NNLO (red). The lower panel shows the perturbative K factors for NLO/LO (turquoise) and NNLO/NLO (mauve).

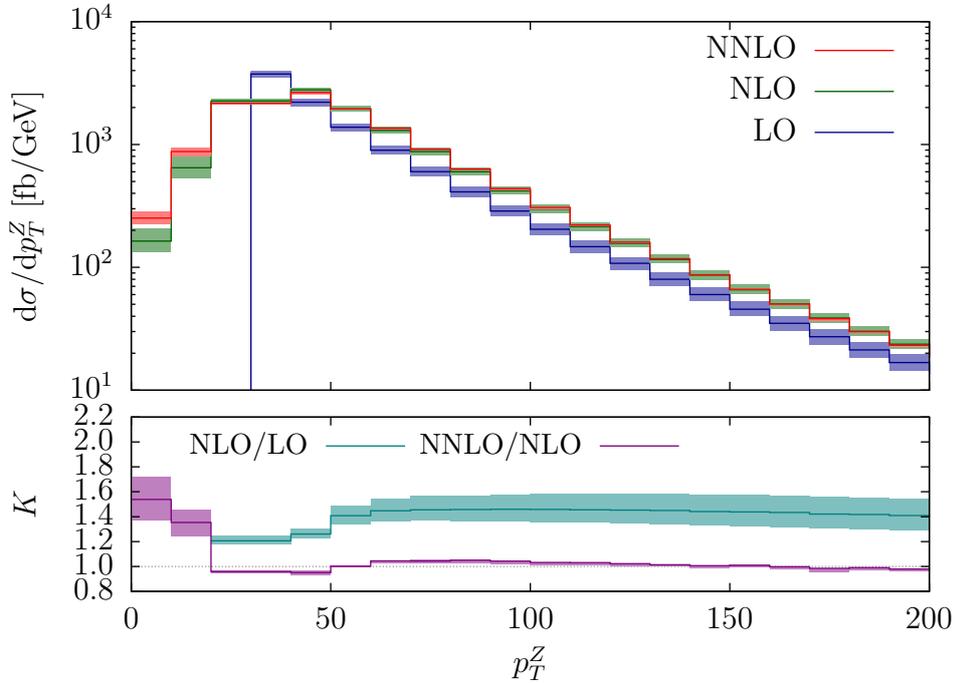


Figure 6.3: The transverse momentum distribution of the Z boson in inclusive Z +jet production in pp collisions with $\sqrt{s} = 8$ TeV at LO (blue), NLO (green), NNLO (red). The lower panel shows the perturbative K factors for NLO/LO (turquoise) and NNLO/NLO (mauve). For the NLO and NNLO predictions, the two bins extending $p_T^Z \in [20, 40]$ GeV have been merged into one.

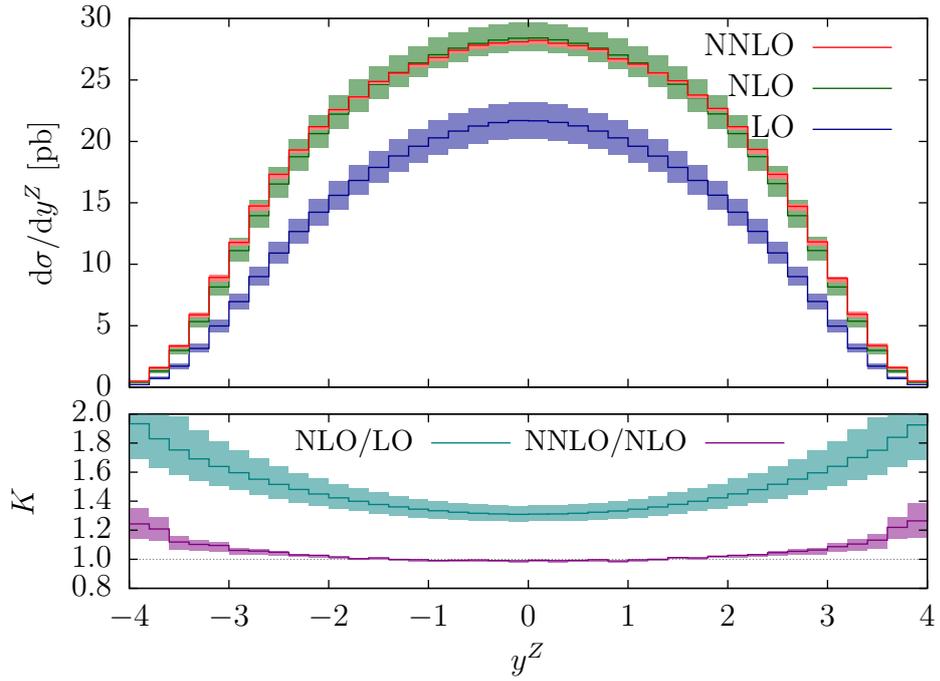


Figure 6.4: The rapidity distribution of the Z boson in inclusive Z +jet production in pp collisions with $\sqrt{s} = 8$ TeV at LO (blue), NLO (green), NNLO (red). The lower panel shows the perturbative K factors for NLO/LO (turquoise) and NNLO/NLO (mauve).

Fig. 6.4 shows the rapidity distribution of the Z boson. The NLO and NNLO corrections are largest in the forward/backward regions where the phase space is enlarged by the possibility that the hadronic radiation partially balances leading to a smaller Z p_T . In these regions, one of the parton momentum fractions is reaching a maximal value. In the central region, the NNLO corrections are very small with a reduced scale dependence.

Fig. 6.5 shows that the NNLO corrections to the transverse momentum distribution for the leptons are uniform and small across the entire range in p_T . Fig. 6.6 shows that the NNLO corrections to the pseudorapidity distribution for the leptons have a clear shape relative to the NLO prediction. In the central regions the NNLO and NLO predictions agree very well, however at large pseudorapidities the NNLO corrections are sizeable with up to 10% corrections over the NLO prediction.

In the differential distributions the corrections are not always uniform, implying that a rescaling of lower order predictions is insufficient for precision applications. The need for using the fully differential higher order predictions can be understood for example in the extraction of parton distributions functions from Z +jet production. At

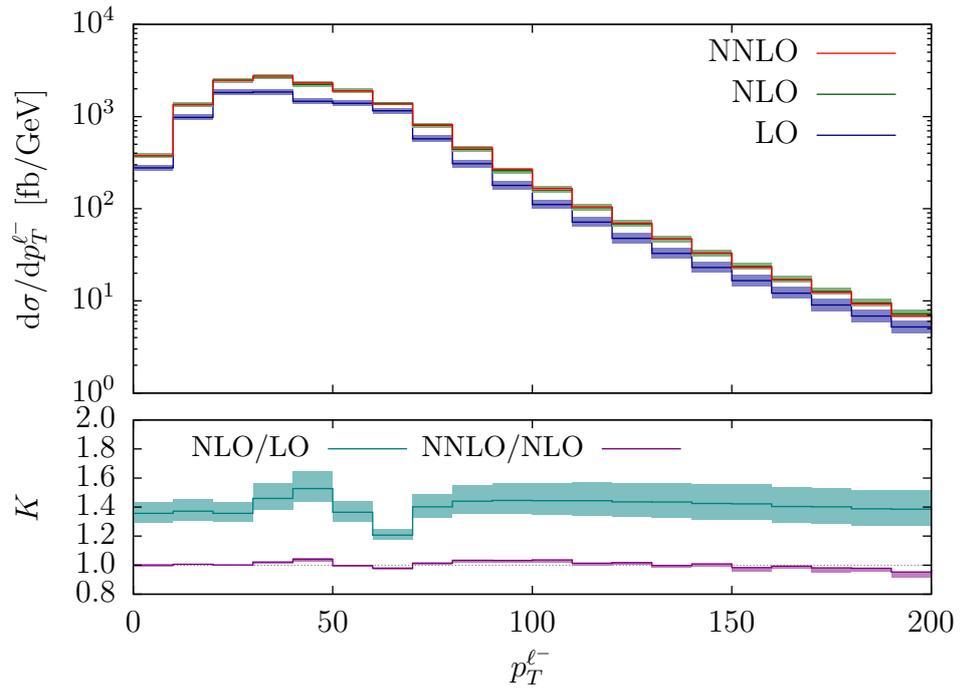


Figure 6.5: The transverse momentum distribution for the negatively charged lepton for Z +jet production in pp collisions with $\sqrt{s} = 8$ TeV at LO (blue), NLO (green), NNLO (red). The lower panel shows the perturbative K factors for NLO/LO (turquoise) and NNLO/NLO (mauve).

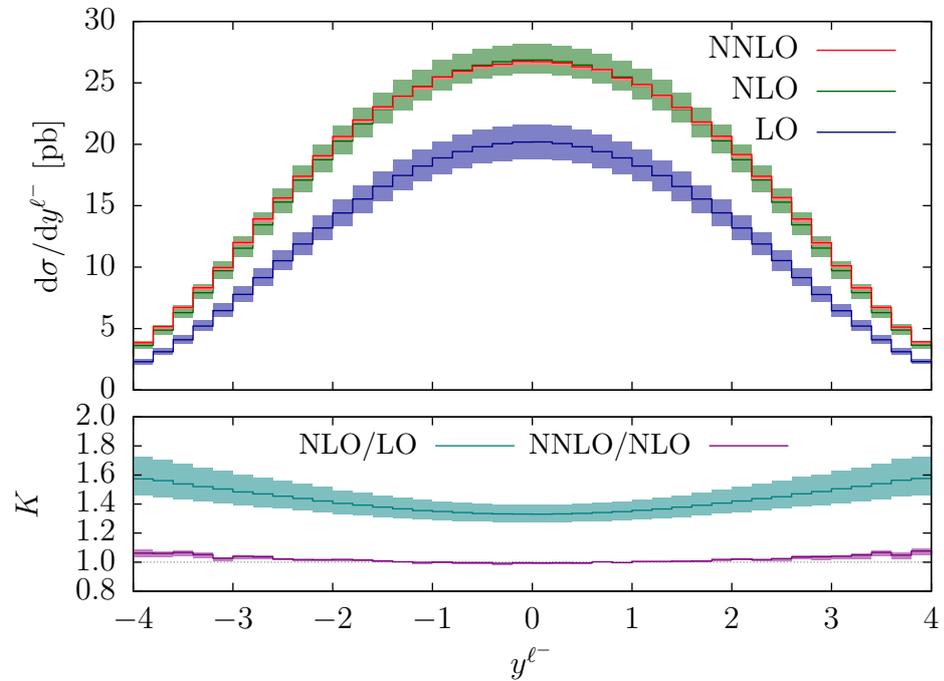


Figure 6.6: The pseudorapidity distribution for the negatively charged lepton for Z +jet production in pp collisions with $\sqrt{s} = 8$ TeV at LO (blue), NLO (green), NNLO (red). The lower panel shows the perturbative K factors for NLO/LO (turquoise) and NNLO/NLO (mauve).

leading order, the momentum fractions of the incoming partons is completely fixed by the transverse momenta and rapidities of the final state particles. At higher orders, the real radiation spoils the leading order kinematics, such that

$$x_1 \geq \frac{1}{\sqrt{s}} \left(\sqrt{(p_T^Z)^2 + m_{\ell\ell}^2} \exp(y_Z) + p_T^{\text{jet}} \exp(y_{\text{jet}}) \right), \quad (6.1.1)$$

$$x_2 \geq \frac{1}{\sqrt{s}} \left(\sqrt{(p_T^Z)^2 + m_{\ell\ell}^2} \exp(-y_Z) + p_T^{\text{jet}} \exp(-y_{\text{jet}}) \right), \quad (6.1.2)$$

where the equality is restored only for the leading order kinematics ($p_T^Z = p_T^{\text{jet}}$). The relevant x ranges probed by Z +jet production is thus determined by the transverse momentum and rapidity distribution of the Z boson and the jet. For our cuts, the smallest momentum fractions probed are $x \sim 8 \cdot 10^{-3}$, and smaller values of x can be attained by enlarging the rapidity interval or by lowering the transverse momentum cut.

6.2 Transverse momentum distribution of the Z boson

For the numerical results in this section, the NNPDF3.0 PDF set [157] is used with the value of $\alpha_s(M_Z) = 0.118$ at NNLO, and $M_Z = 91.1876$ GeV. As in the previous section, the same PDF choice is used for the NLO and NNLO prediction, using the same choice of $\alpha_s(M_Z)$ throughout. The factorisation and renormalisation scales are chosen dynamically on an event-by-event basis as,

$$\mu \equiv \mu_R = \mu_F = \sqrt{m_{\ell\ell}^2 + (p_T^Z)^2}, \quad (6.2.1)$$

where $m_{\ell\ell}$ and p_T^Z are the invariant mass and the transverse momentum of the final state lepton pair respectively. The theoretical uncertainty is estimated by varying the scale choice by a factor in the range $[1/2, 2]$. A cut on the transverse momentum of the Z boson, $p_T^Z > 20$ GeV, is also included. In the low transverse momentum region, large logarithmic corrections of the form $\log^n(p_T/M_Z)$ appear at each order in the perturbative expansion in α_s , spoiling its convergence. A reliable theoretical prediction in this region can only be obtained by resummation [119, 120, 158, 159] of these logarithms to all orders in perturbation theory. In the computation of the inclusive lepton pair

| | ATLAS | CMS |
|-------------------|--|--|
| leading lepton | $ \eta_{\ell_1} < 2.4$ $p_T^{\ell_1} > 20$ GeV | $ \eta_{\ell_1} < 2.1$ $p_T^{\ell_1} > 25$ GeV |
| subleading lepton | $ \eta_{\ell_2} < 2.4$ $p_T^{\ell_2} > 20$ GeV | $ \eta_{\ell_2} < 2.4$ $p_{T,2}^{\ell_2} > 10$ GeV |

Table 6.2: Kinematical cuts used to define the fiducial phase space for the leading and subleading leptons, ordered in transverse momentum, for the measurements of ATLAS [143] and CMS [145].

production cross section used to normalise the transverse momentum distributions, the same scale (6.2.1) is used but varied independently over the same range of scale variation. The inclusive cross section in this case is, however, dominated by the regime where $\mu^2 \approx m_{\ell\ell}^2$.

The experimental measurements of Z boson production at finite transverse momentum are presented in the form of fiducial cross sections over a restricted phase space for the final state leptons, which is fully contained in the detector's coverage. The NNLO corrections to the transverse momentum distribution can be compared to data by considering the same cuts to the lepton kinematics as presented in the ATLAS [143] and CMS [145] analyses using data from Run 1 of the LHC with $\sqrt{s} = 8$ TeV, which are summarised in Tab. 6.2. In this section we will focus on the ATLAS measurement in the fiducial region defined by a broad dilepton invariant mass window around the Z resonance, $66 \text{ GeV} < m_{\ell\ell} < 116 \text{ GeV}$, and compare both the absolute and normalised p_T^Z distribution to the experimental data.

Fig. 6.7 shows a comparison of data from the ATLAS analysis within this fiducial region. At low transverse momentum, the NNLO correction increases the cross section by about 6% (compared to NLO) and significantly reduces the scale uncertainty. However, there is still significant tension between the ATLAS data and the NNLO prediction.

The data presented in Fig. 6.7 does not include the error on the integrated luminosity, which amounts to an overall normalisation uncertainty of 2.8% on all data points. To cancel the luminosity uncertainty from the measured data, it is more appropriate to normalise the data by the measured values for the inclusive lepton pair cross section

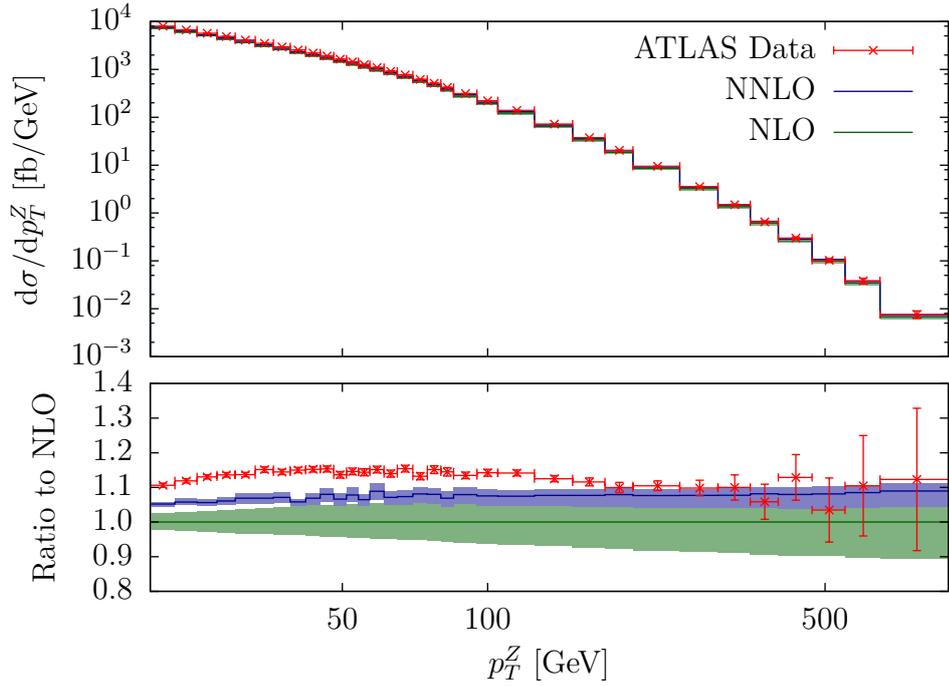


Figure 6.7: The unnormalised Z boson transverse momentum distribution for the cuts given in Tab. 6.2 and $66 \text{ GeV} < m_{\ell\ell} < 116 \text{ GeV}$. ATLAS data is taken from Ref. [143]. The luminosity error is not shown. The green bands denote the NLO prediction with scale uncertainty and the blue bands show the NNLO prediction with scale uncertainty.

in this fiducial bin. The cross section for this mass window was measured to be [143],

$$\sigma_{\text{exp}}(66 \text{ GeV} < m_{\ell\ell} < 116 \text{ GeV}) = 537.10 \pm 0.45\% \text{ (sys.)} \pm 2.80\% \text{ (lumi.) pb.}$$

For the NNLO prediction of the cross section in this fiducial region, we obtain,

$$\sigma_{\text{NNLO}}(66 \text{ GeV} < m_{\ell\ell} < 116 \text{ GeV}) = 507.9_{-0.7}^{+2.4} \text{ pb.}$$

We see that there is some tension between the measured cross section compared to the theoretical result. The normalised distribution is presented in Fig. 6.8, where excellent agreement is observed between the normalised NNLO prediction and the experimental data across a wide range of transverse momentum. The scale bands on the normalised theory predictions are obtained by independently varying the scale in the numerator and denominator, where we use the NNLO prediction for the inclusive Drell–Yan cross section in the normalisation throughout.

The EW corrections to this distribution are known to be small at moderate transverse momenta but become more sizeable in the high- p_T^Z tail where they can reach

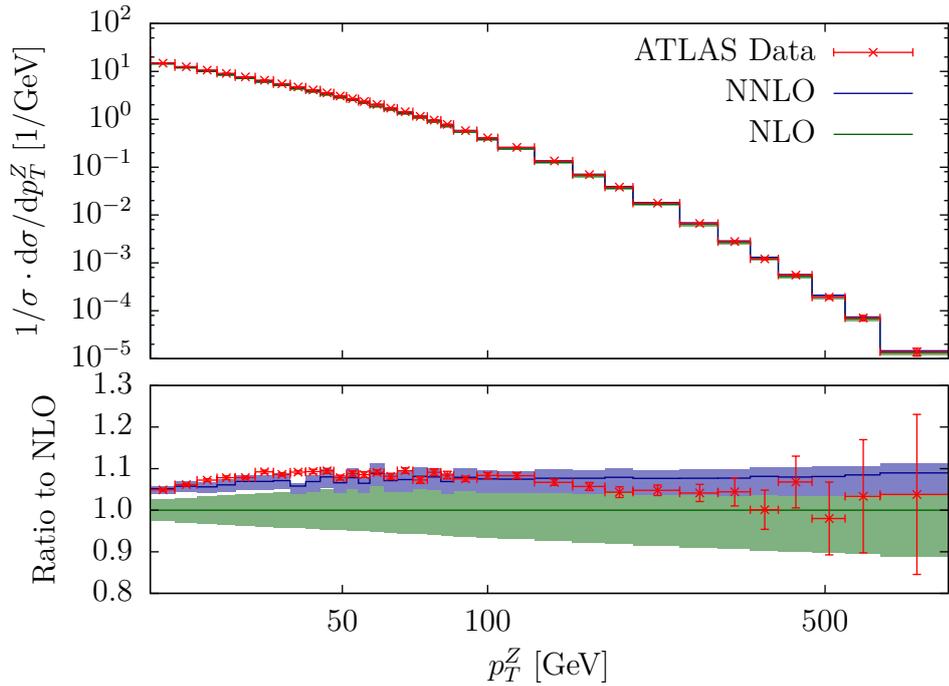


Figure 6.8: The normalised Z boson transverse momentum distribution for the cuts given in Tab. 6.2 and $66 \text{ GeV} < m_{\ell\ell} < 116 \text{ GeV}$. ATLAS data is taken from Ref. [143]. The green bands denote the NLO prediction with scale uncertainty and the blue bands show the NNLO prediction with scale uncertainty.

$\sim -15\%$ for $p_T^Z \gtrsim 600 \text{ GeV}$ [128]. It can be therefore expected that the inclusion of EW corrections will further improve the prediction of the shape in the tail. However, this region of the distribution is still dominated by the statistical uncertainties of the experimental data and a fully consistent inclusion of EW corrections is beyond the scope of this work.

6.2.1 Double-differential distributions

We now compare our theoretical predictions for dilepton production at large transverse momentum for different ranges of (a) the dilepton mass and (b) the dilepton rapidity. These double-differential distributions have also been studied by ATLAS [143] and CMS [145] with Run 1 data using the lepton rapidity and transverse momentum cuts summarised in Tab. 6.2.

We first consider the ATLAS measurement, which covers a broader kinematical range. In Fig. 6.9 we present the unnormalised double-differential distribution with respect to the transverse momentum of the Z boson and the invariant mass of the lepton pair, $m_{\ell\ell}$, normalised to the NLO prediction and compare it to the ATLAS

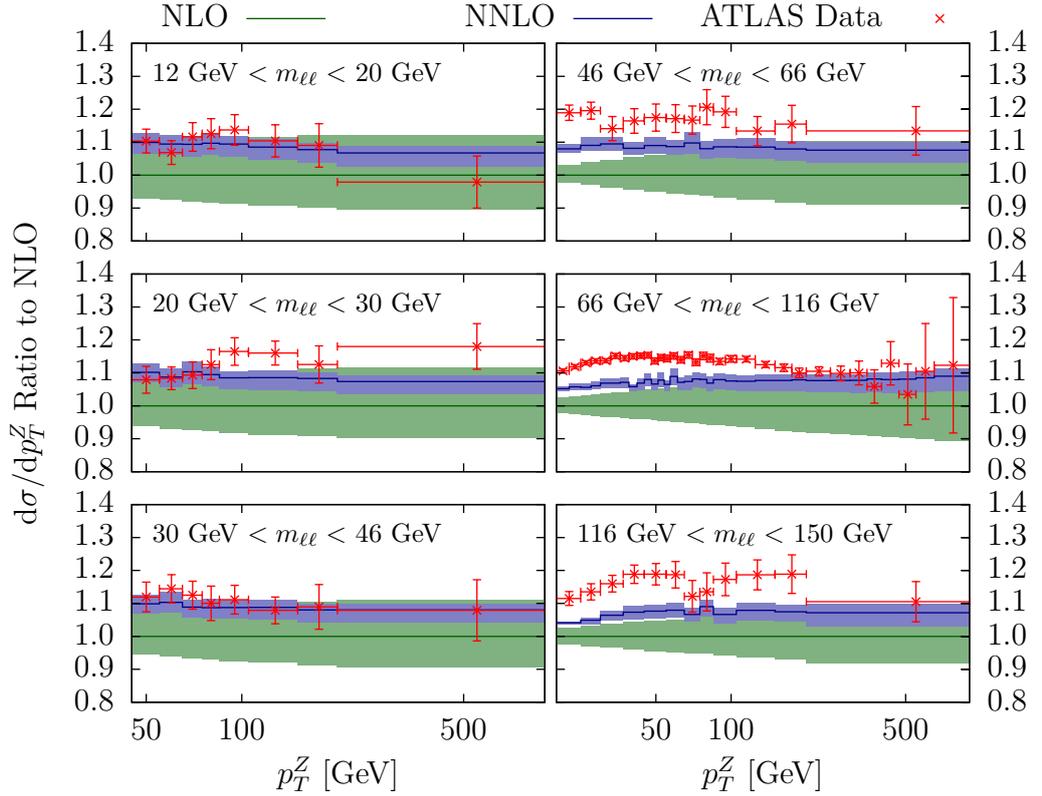


Figure 6.9: The unnormalised double-differential transverse momentum distribution for the Z boson in windows of invariant mass of the leptons, $m_{\ell\ell}$, with a rapidity cut on the Z boson of $|y^Z| < 2.4$. The ATLAS data is taken from Ref. [143]. The luminosity error is not shown. The green bands denote the NLO prediction with scale uncertainty and the blue bands show the NNLO prediction with scale uncertainty.

data from Ref. [143]. We observe that there is good agreement between the NNLO prediction and the data in the low $m_{\ell\ell}$ mass windows. For values of $m_{\ell\ell}$ close to the Z boson mass, where the statistical accuracy of the data is highest, the NNLO prediction is below the data by about 5-8%.

To obtain normalised distributions, these data sets are divided by the inclusive dilepton cross sections for each fiducial bin, defined by the lepton cuts given in Tab. 6.2 and the appropriate dilepton invariant mass cut. We computed these fiducial cross sections to NNLO accuracy ($\mathcal{O}(\alpha_s^2)$) and compare them to the measured values quoted by ATLAS [143] in Fig. 6.10. We clearly observe that the central value of the NNLO prediction falls below the experimental data in all mass bins.

In the three lowest mass bins, the experimental analysis performed by ATLAS [143] includes an additional kinematical cut of $p_T^Z > 45$ GeV, in addition to the fiducial region defined in Tab. 6.2. This means that the normalisation choice defined in Eq. (3.4.2)

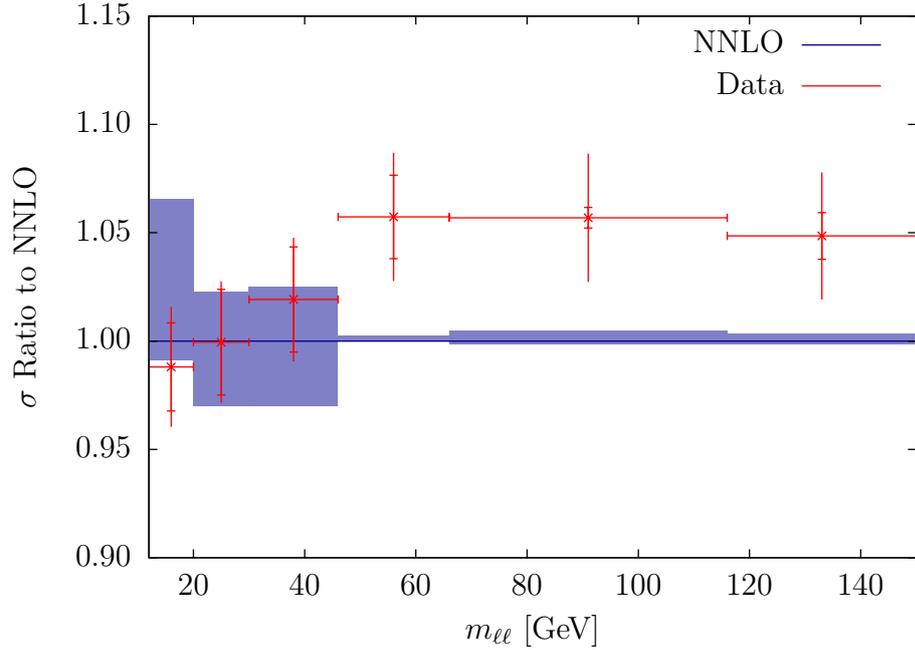


Figure 6.10: The inclusive dilepton cross section for the same $m_{\ell\ell}$ bins as in Fig. 6.9 and with a rapidity cut on the Z boson of $|y^Z| < 2.4$. The experimental data is taken from the ATLAS analysis in Ref. [143]. The ticks on the vertical error bands denote the systematic uncertainty from the measurement, the vertical bars without the ticks are the luminosity uncertainty only. The blue bands show the NNLO prediction with scale uncertainty. The theoretical cross section predictions for the invariant mass windows $12 < m_{\ell\ell} < 20$ GeV, $20 < m_{\ell\ell} < 30$ GeV, and $30 < m_{\ell\ell} < 46$ GeV are calculated using the Z +jet calculation since the additional kinematical cut $p_T^Z > 45$ GeV prevents contributions from the LO Drell–Yan configuration in these invariant mass windows.

is only NLO accurate in this fiducial region. Indeed, the normalisation choice can be promoted to NNLO accuracy by normalising to the full Z +jet calculation,

$$\sigma\Big|_{p_T^Z > 45 \text{ GeV}} = \int_{45 \text{ GeV}}^{\infty} \frac{d\sigma}{dp_T^Z} dp_T^Z = \left(\sigma_{\text{LO}}^{\text{Z+jet}} + \sigma_{\text{NLO}}^{\text{Z+jet}} + \sigma_{\text{NNLO}}^{\text{Z+jet}} \right)\Big|_{p_T^Z > 45 \text{ GeV}} + \mathcal{O}(\alpha_s^4). \quad (6.2.2)$$

For the remainder of this section, this choice of normalisation will be used for the three lower mass bins to give the most accurate result possible. In the three bins with larger $m_{\ell\ell}$, the scale uncertainty on the NNLO prediction is below 0.7%, which results in tension between data and theory at the level of two standard deviations. Combining together the unnormalised differential distribution with the inclusive cross sections, we obtain the normalised distributions shown in Fig. 6.11. For all of the $m_{\ell\ell}$ windows the tension between the data and NNLO theory is largely relieved. At the highest values of p_T^Z , the tendency of the data to fall below the theory prediction may be an

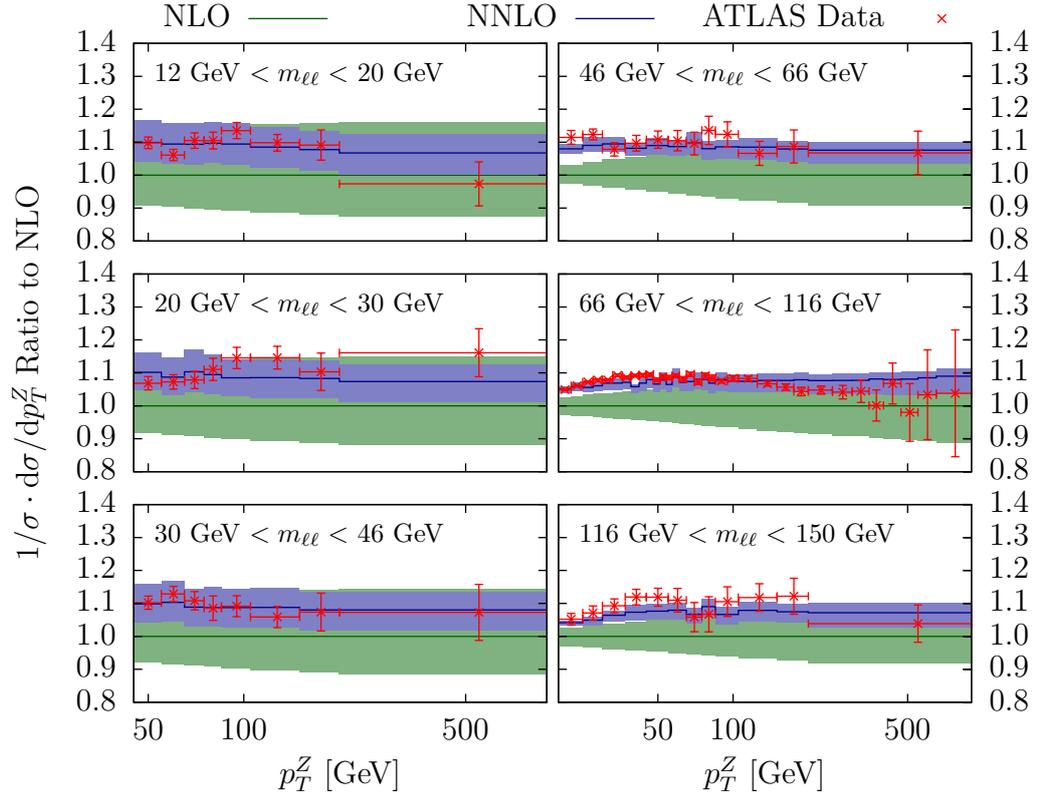


Figure 6.11: The normalised double-differential transverse momentum distribution for the Z boson in windows of invariant mass of the leptons, $m_{\ell\ell}$, with a rapidity cut on the Z boson of $|y^Z| < 2.4$. The ATLAS data is taken from Ref. [143]. The green bands denote the NLO prediction with scale uncertainty and the blue bands show the NNLO prediction with scale uncertainty. The normalisation cross section for the invariant mass windows $12 < m_{\ell\ell} < 20$ GeV, $20 < m_{\ell\ell} < 30$ GeV, and $30 < m_{\ell\ell} < 46$ GeV are calculated using the Z +jet calculation since the kinematical cut $p_T^Z > 45$ GeV prevents contributions from the LO Drell–Yan configuration in these invariant mass windows.

indication of the onset of EW corrections [128], which are negative in this region. Any remaining tension for medium values of p_T^Z could potentially be accounted for revisiting the parton distribution functions (especially the gluon distribution) in the kinematical region relevant to this measurement.

The same tension between NNLO theory and ATLAS data for the unnormalised distribution is visible in Fig. 6.12, which shows the unnormalised double-differential distribution with respect to the transverse momentum of the Z boson for 66 GeV $< m_{\ell\ell} < 116$ GeV in different ranges of the rapidity of the Z boson, normalised to the NLO prediction. The NNLO corrections are uniform in rapidity and transverse momentum at the level of about 5–8%, and they decrease the residual theoretical uncertainty to 2–4%. The data, which correspond to the invariant mass bin containing the Z resonance,

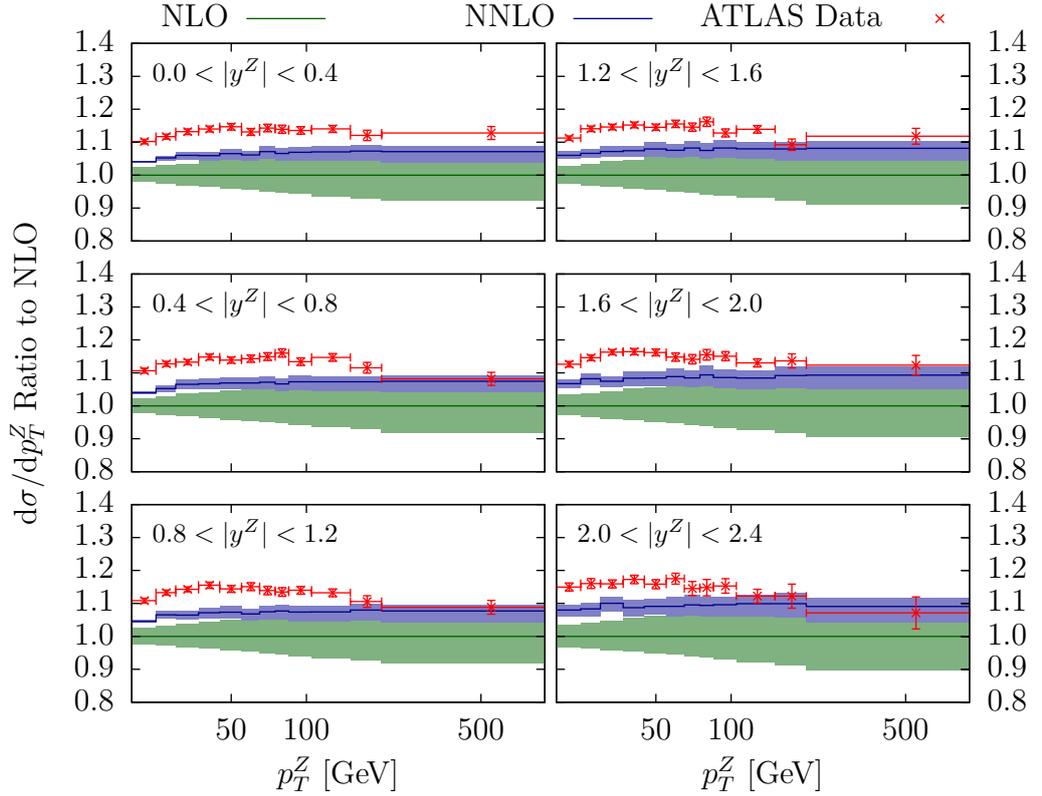


Figure 6.12: The unnormalised double-differential transverse momentum distribution for the Z boson in windows of rapidity of the Z boson, y^Z , with an invariant mass cut on the final state leptons of $66 \text{ GeV} < m_{\ell\ell} < 116 \text{ GeV}$. The ATLAS data is taken from Ref. [143]. The luminosity error is not shown. The green bands denote the NLO prediction with scale uncertainty and the blue bands show the NNLO prediction.

are consistently above the NNLO prediction.

In order to eliminate the overall luminosity uncertainty, one again normalises the data to the inclusive fiducial dilepton cross section in the respective bins in y^Z . Fig. 6.13 shows the ratio of ATLAS data and the NNLO cross section compared to the NLO prediction¹. As for the rapidity-integrated cross section in this mass bin, Fig. 6.10, we observe the NNLO prediction to be accurate to 0.9%, and to underestimate the ATLAS data by about two standard deviations.

The normalised double-differential distribution is shown in Fig. 6.14. There is excellent agreement between the normalised NNLO prediction and the ATLAS data for all rapidity bins, again offering the possibility of further constraining the PDFs in

¹Note that the ATLAS data is extracted from Ref. [143] by summing the bins of the p_T^Z distributions in the various y^Z bins. We have checked that this procedure applied to the different $m_{\ell\ell}$ bins shown in Fig. 6.10 correctly reproduces the data.

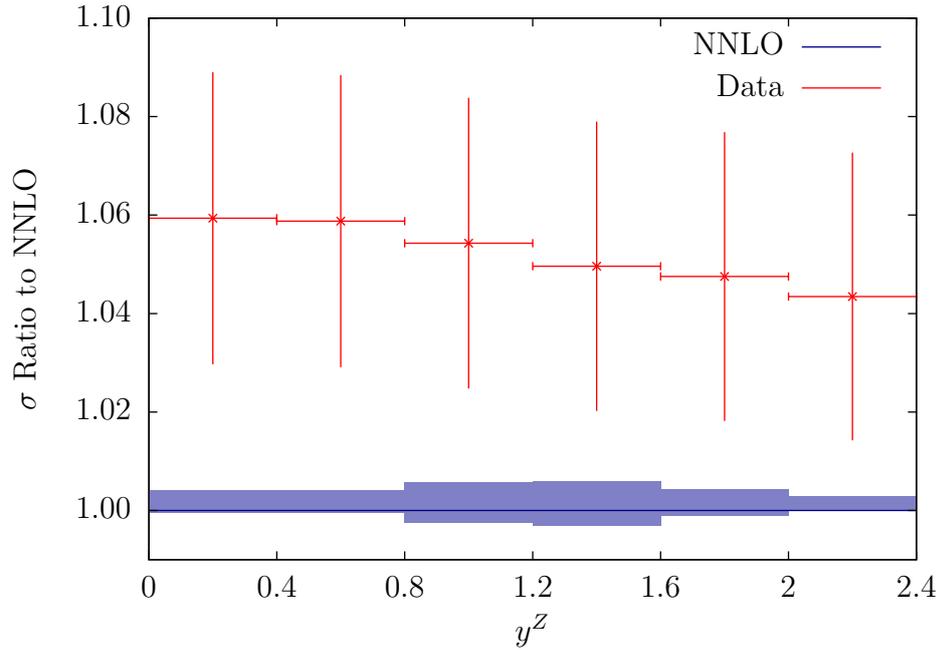


Figure 6.13: The inclusive dilepton cross section for the same $|y^Z|$ bins as in Fig. 6.12 and an invariant mass cut on final state leptons of $66 \text{ GeV} < m_{\ell\ell} < 116 \text{ GeV}$. The ATLAS data is extracted from Ref. [143] by summing up the transverse momentum distributions in the respective bins. The vertical error bars are given by the luminosity uncertainty. The blue bands show the NNLO prediction with scale uncertainty.

this kinematic region.

The CMS measurement of the Z boson transverse momentum distribution at 8 TeV [145] concentrates on the dilepton invariant mass range $81 \text{ GeV} < m_{\ell\ell} < 101 \text{ GeV}$ around the Z resonance. The fiducial region of this measurement, defined in Tab. 6.2, applies asymmetric cuts on the leading and subleading lepton, disregarding the lepton charge. To ensure the correct implementation of these cuts in NNLOJET, and to maintain local cancellations between matrix elements and subtraction terms in the antenna subtraction method, the lepton identification as leading and subleading had to be passed alongside the lepton charge through all subtraction terms and into the event analysis routines. The double-differential measurement is performed in p_T^Z and y^Z . Comparing the absolute double-differential distributions of Fig. 6.15 to our NNLO predictions, we observe the same features as for the ATLAS measurement, with positive NNLO corrections at the level of 5–8% and an NNLO scale uncertainty of 2–4%. Compared to NLO, inclusion of NNLO corrections brings the theoretical prediction closer to the experimental data, which are however still about 5–8% larger than expected from

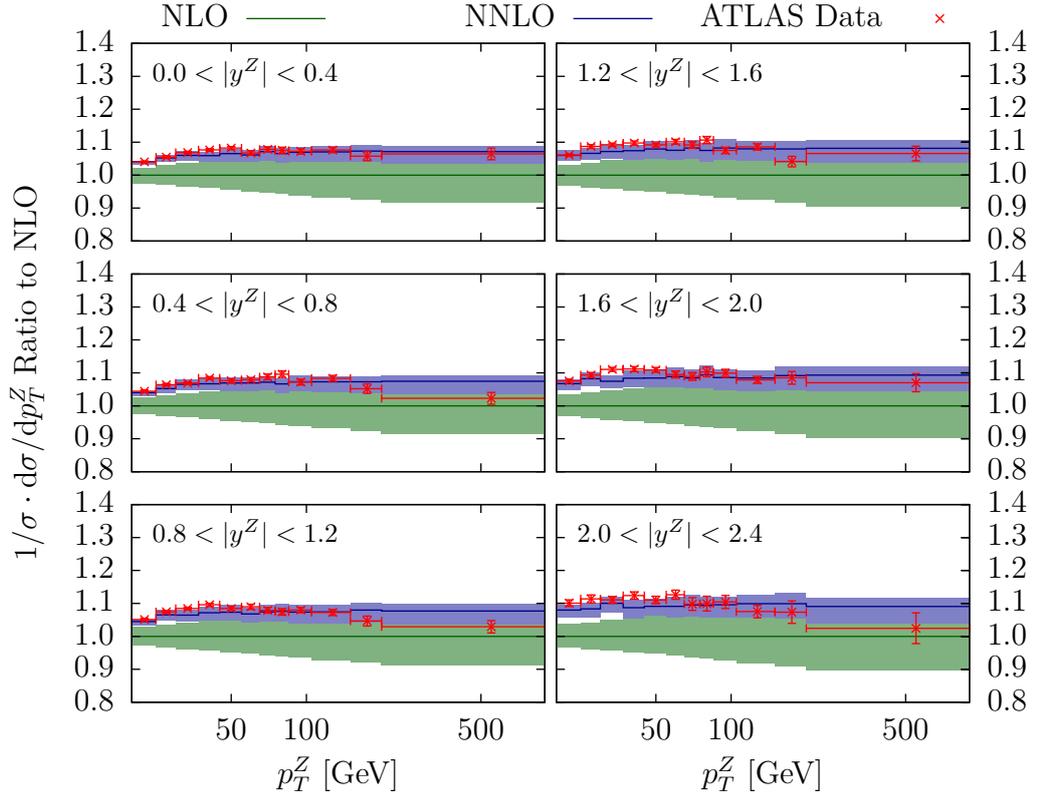


Figure 6.14: The normalised double-differential transverse momentum distribution for the Z boson in windows of rapidity of the Z boson, y^Z , with an invariant mass cut on final state leptons of $66 \text{ GeV} < m_{\ell\ell} < 116 \text{ GeV}$. The ATLAS data is taken from Ref. [143]. The green bands denote the NLO prediction with scale uncertainty and the blue bands show the NNLO prediction.

theory.

In the CMS study [145], the normalised cross section is obtained by dividing the distributions by the fiducial cross section integrated over all bins in y^Z (in contrast to ATLAS, which normalised to the fiducial cross sections only integrated over the respective bin in y^Z). At NNLO, we obtain for the fiducial cross section with CMS cuts:

$$\sigma_{\text{NNLO}}(81 \text{ GeV} < m_{\ell\ell} < 101 \text{ GeV}) = 450.6_{-1.6}^{+2.7} \text{ pb.}$$

The normalised distributions of CMS are compared to theory in Fig. 6.16, where excellent agreement is observed upon inclusion of the NNLO corrections.

Compared to NNLO theory, both ATLAS and CMS fiducial cross section measurements of the Z boson transverse momentum display a similar pattern of disagreement for the absolute distributions while having excellent agreement for the normalised dis-

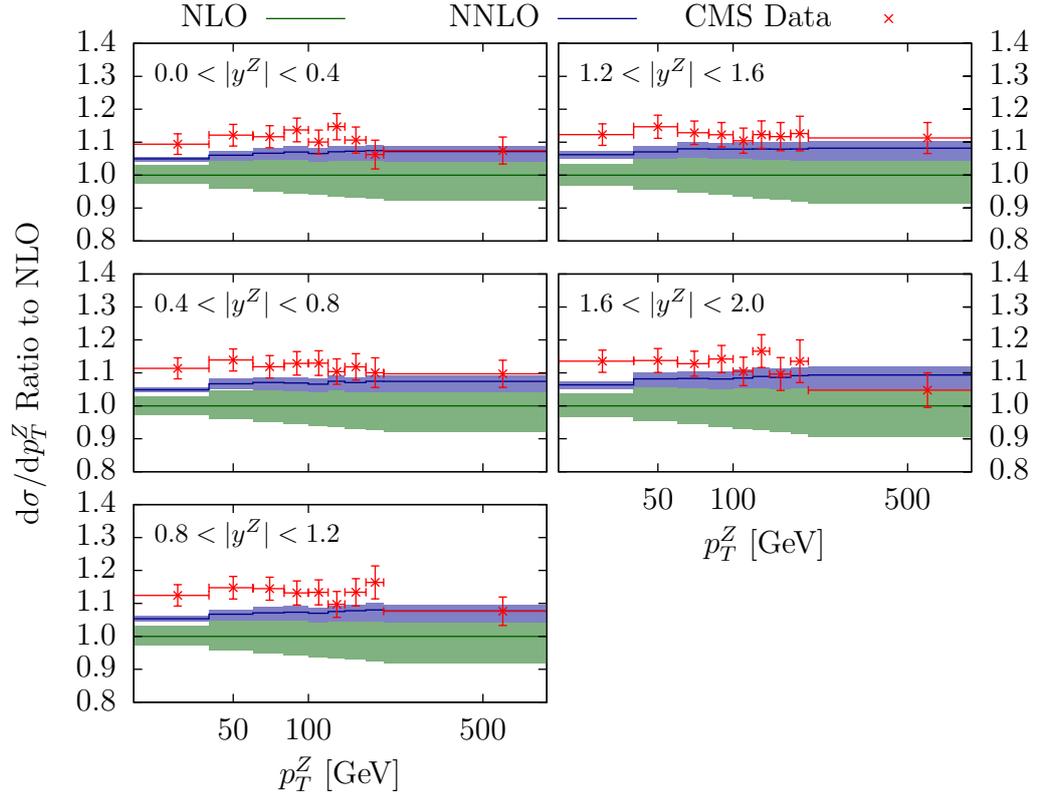


Figure 6.15: The unnormalised double-differential transverse momentum distribution for the Z boson in windows of rapidity of the Z boson, y^Z , with an invariant mass cut on final state leptons of $81 \text{ GeV} < m_{\ell\ell} < 101 \text{ GeV}$. The CMS data is taken from Ref. [145]. The experimental error bars include the luminosity uncertainty. The green bands denote the NLO prediction with scale uncertainty and the blue bands show the NNLO prediction.

tributions. The inclusion of the newly derived NNLO corrections to the transverse momentum distribution are crucial for a meaningful comparison between data and theory: (a) they reduce the theory uncertainty to a level that firmly establishes the discrepancy on the absolute cross sections, and (b) they modify the central value and the shape of the theory prediction to better agree with the data on the normalised distributions.

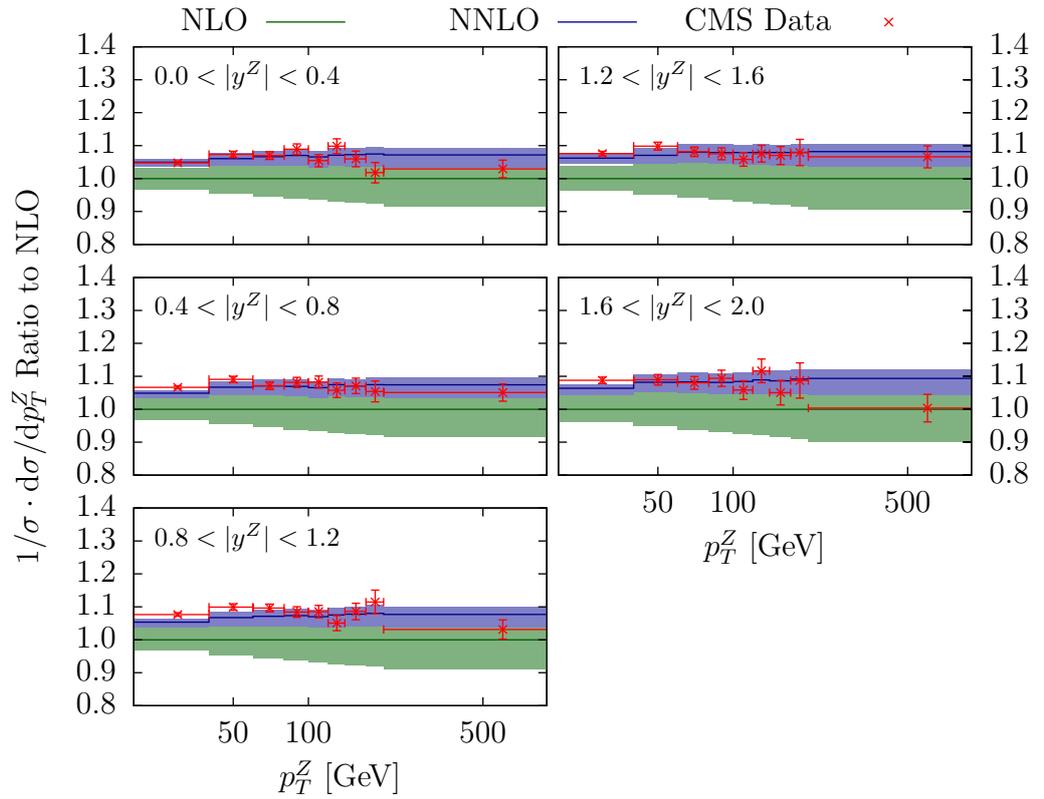


Figure 6.16: The normalised double-differential transverse momentum distribution for the Z boson in windows of rapidity of the Z boson, y^Z , with an invariant mass cut on final state leptons of $81 \text{ GeV} < m_{\ell\ell} < 101 \text{ GeV}$. The CMS data is taken from Ref. [145]. The green bands denote the NLO prediction with scale uncertainty and the blue bands show the NNLO prediction.

Chapter 7

Conclusions

The research within this thesis has focused on using the antenna subtraction method to calculate the NNLO corrections to a non-trivial LHC process: Z bosons in association with a hadronic jet. This is the first example of the antenna subtraction used to perform a calculation including *all* initial states at all colour contributions for a LHC process with jets in the final state.

To perform this calculation, it was clear that it was highly desirable to construct a flexible, extensible code that can be used to calculate arbitrarily complicated final states relevant for LHC processes. Like any complicated Monte Carlo code, significant numerical optimisations are required to perform calculations within a reasonable time scale with the computing resources available. The optimisations discussed in chapter 4 are now readily available for all current and future projects using the antenna subtraction method. Any future optimisations implemented can be immediately applied to the Z +jet calculation through the use of the auto-generated subroutines.

Much of the previous work within the NNLOJET collaboration has focused on the construction of the subtraction terms. With the completion of the Z +jet [1] and H +jet [160] calculations, all of the subtraction terms relevant for one jet final states are now complete, meaning that any process with a single hadronic jet in the final state can be calculated, provided that all of the matrix elements relevant for the NNLO corrections are available. Indeed, it is clear now that the primary limitation in performing NNLO calculations is the lack of two-loop matrix elements available. Whereas one-loop matrix elements have largely been automated to extremely high parton multiplicities,

two-loop matrix elements still pose a significant computational challenge. It is this limitation that severely undermines any motivation for a fully automated approach to NNLO subtraction with large parton multiplicities.

It is clear then that the future of the NNLOJET development should focus on implementing processes involving either a colourless final state or a one-jet final state, given that these subtraction terms are now fully under control. In the short term, a calculation of W^\pm +jet is straightforward using the framework established in this thesis. The W matrix elements are a simplification of the equivalent Z boson matrix elements, which can be implemented with relative ease. This will be important to validate the conclusions of Ref. [76].

Beyond this, the next natural step is to include processes with a more complicated colourless final state, for example Higgs boson production in association with a Z boson. Here the matrix elements are all known or can be calculated from the pure Z boson matrix elements. The challenge that arises with this sort of process is the increased dimensionality of the phase space which will prove to be even more of a challenge numerically.

The phenomenological results of the Z +jet is one of several calculations in a new era of precision LHC physics. Scale uncertainties are of the order of one percent with a process that is experimentally very clean. This is a powerful test of the Standard Model on the boundary between EW and QCD physics.

One of the main motivations for calculating the NNLO corrections to Z +jet was to use the calculation to constrain the gluon PDF. To perform a PDF fit requires the calculation to be repeated for a huge number of potential PDF sets to tune the calculation to data. Given the severe numerical cost of performing the calculation, it is clear that this is not feasible by a brute force approach.

Useful tools for PDF fitters are `applgrid` [153] and `fastNLO` [154], which create lookup tables for partonic cross section coefficients for each observable of interest. The value of this is that the PDF set can be changed *a posteriori*, meaning that the numerically expensive calculation does not need to be repeated for each PDF set. Work is on-going to construct an interface for both `applgrid` and `fastNLO` to the NNLOJET framework which, once completed, will be the first NNLO accurate calculation interfaced with their respective codes. This is not a trivial task, the statistical challenges

encountered in chapter 5 are particularly prevalent because of the double unresolved configuration of the double real phase space. Precisely how to combine multiple grids to smooth out outliers without artificially biasing the results will need to be resolved in due course. The performance impact during the filling of the grids will also need to be assessed.

Using a similar line of argument, it would also be advantageous to use the n -tuples framework [155] for the Z +jet NNLO calculation. n -tuples have the added advantage compared to applgrid and fastNLO that final state cuts can also be changed *a posteriori*. The cost of this flexibility is that the n -tuple files are significantly larger than the grid files from applgrid and fastNLO. The first attempts to extend the n -tuples framework to NNLO have been presented for the calculation $e^+e^- \rightarrow 3$ jets [161, 162]. In principle, this framework could be used for NNLO calculations within a hadron collider environment, however the impact on calculation time and storage size would need to be assessed.

The calculation of the NNLO corrections to the Z boson transverse momentum distribution firmly establishes that the discrepancy between theory and data lies in the total cross section. Figs. 6.10 and 6.13 demonstrate that there is tension between the 8 TeV ATLAS data and the NNLO prediction using the NNPDF3.0 PDF set. Similar conclusions can be drawn from the CMS results.

Once the distributions have been normalised by their total cross section, the agreement between the NNLO prediction and the experimental data is excellent across a wide range of transverse momentum, both in the single differential distribution and the double-differential results with respect to $m_{\ell\ell}$ and $|y^Z|$. This is a huge success for Standard Model phenomenology, where both the experimental and theoretical uncertainties are small and are in agreement.

The calculation presented warrants a thorough investigation of all of the contributing parameters to assess the impact of the tension in the total cross section between the theoretical prediction and the experimental data. The calculation should be repeated using many different PDF sets and include PDF uncertainties, lending itself to an interface with applgrid and fastNLO. The impact of the variation of the strong coupling constant should also be assessed in a similar fashion.

Another extension would be to include the NLO EW corrections from Refs. [128,

129]. It was argued that in the context of the experimental uncertainties that these corrections should be small. With 13 TeV data from Run 2 of the LHC on the horizon, this may no longer be the case. Indeed, a 13 TeV prediction for the transverse momentum distribution, including all calculated theoretical corrections, will be an extremely valuable extension to the work presented.

The transverse momentum distribution was presented with a large cut, $p_T^Z > 20$ GeV, to maintain the validity of the perturbative expansion of α_s . To create a complete description of the transverse momentum distribution, the calculation should be matched to the state-of-the-art resummation calculations for this observable [119, 120]. Indeed, it could be of great interest to drive the transverse momentum cut as low as possible to match logarithmic coefficients to a resummed calculation. Lowering the p_T^Z cut will have a severe impact on the numerical stability of the code however.

The research presented within this thesis opens up new opportunities for the field of precision phenomenology. Whereas once NNLO corrections were limited to a few of the partonic channels for a handful of processes, there are now many full NNLO calculations readily available. The analytical challenges of NNLO subtraction for low parton multiplicities have been resolved, leaving only the question of whether a numerical implementation is feasible. The NNLOJET programme will provide a large array of processes within a consistent framework, exploiting the many numerical efficiency upgrades across all of the processes within the package. This will be extremely valuable for the wider theoretical and experimental communities to perform a wide range of analyses and comparisons to data in the years to come.

Bibliography

- [1] A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss, and T. A. Morgan. “Precise QCD predictions for the production of a Z boson in association with a hadronic jet”. *Phys. Rev. Lett.* 117.2 (2016), p. 022001. arXiv: 1507.02850 [hep-ph].
- [2] A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss, and T. A. Morgan. “The NNLO QCD corrections to Z boson production at large transverse momentum”. *JHEP* 07 (2016), p. 133. arXiv: 1605.04295 [hep-ph].
- [3] A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss, and T. A. Morgan. “Z+jet production at NNLO”. *13th DESY Workshop on Elementary Particle Physics: Loops and Legs in Quantum Field Theory (LL2016) Leipzig, Germany, April 24-29, 2016*. 2016. arXiv: 1607.01749 [hep-ph].
- [4] A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, A. Huss, and T. A. Morgan. “NNLO QCD corrections for Z boson plus jet production”. *Proceedings, 12th International Symposium on Radiative Corrections (Radcor 2015) and LoopFest XIV (Radiative Corrections for the LHC and Future Colliders)*. 2016. arXiv: 1601.04569 [hep-ph].
- [5] ATLAS Collaboration. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. *Phys. Lett.* B716 (2012), pp. 1–29. arXiv: 1207.7214 [hep-ex].
- [6] CMS Collaboration. “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC”. *Phys. Lett.* B716 (2012), pp. 30–61. arXiv: 1207.7235 [hep-ex].

- [7] J. R. Andersen et al. “Les Houches 2013: Physics at TeV Colliders: Standard Model Working Group Report” (2014). arXiv: 1405.1067 [hep-ph].
- [8] F. Bloch and A. Nordsieck. “Note on the Radiation Field of the electron”. *Phys. Rev.* 52 (1937), pp. 54–59.
- [9] T. Kinoshita. “Mass singularities of Feynman amplitudes”. *J. Math. Phys.* 3 (1962), pp. 650–677.
- [10] T. D. Lee and M. Nauenberg. “Degenerate Systems and Mass Singularities”. *Phys. Rev.* 133 (6B 1964), B1549–B1562.
- [11] A. Gehrmann-De Ridder, T. Gehrmann, and E. W. N. Glover. “Antenna subtraction at NNLO”. *JHEP* 09 (2005), p. 056. arXiv: hep-ph/0505111 [hep-ph].
- [12] A. Gehrmann-De Ridder, T. Gehrmann, and E. W. N. Glover. “Infrared structure of $e^+e^- \rightarrow 2$ jets at NNLO”. *Nucl. Phys.* B691 (2004), pp. 195–222. arXiv: hep-ph/0403057 [hep-ph].
- [13] A. Gehrmann-De Ridder, T. Gehrmann, and E. W. N. Glover. “Gluon-gluon antenna functions from Higgs boson decay”. *Phys. Lett.* B612 (2005), pp. 49–60. arXiv: hep-ph/0502110 [hep-ph].
- [14] A. Gehrmann-De Ridder, T. Gehrmann, and E. W. N. Glover. “Quark-gluon antenna functions from neutralino decay”. *Phys. Lett.* B612 (2005), pp. 36–48. arXiv: hep-ph/0501291 [hep-ph].
- [15] S. L. Glashow. “Partial Symmetries of Weak Interactions”. *Nucl. Phys.* 22 (1961), pp. 579–588.
- [16] S. Weinberg. “A Model of Leptons”. *Phys. Rev. Lett.* 19 (1967), pp. 1264–1266.
- [17] A. Salam. “Weak and Electromagnetic Interactions”. *Conf. Proc.* C680519 (1968), pp. 367–377.
- [18] F. Englert and R. Brout. “Broken Symmetry and the Mass of Gauge Vector Mesons”. *Phys. Rev. Lett.* 13 (1964), pp. 321–323.
- [19] P. W. Higgs. “Broken Symmetries and the Masses of Gauge Bosons”. *Phys. Rev. Lett.* 13 (1964), pp. 508–509.

- [20] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble. “Global Conservation Laws and Massless Particles”. *Phys. Rev. Lett.* 13 (1964), pp. 585–587.
- [21] L. D. Faddeev and V. N. Popov. “Feynman Diagrams for the Yang-Mills Field”. *Phys. Lett.* B25 (1967), pp. 29–30.
- [22] M. E. Peskin and D. V. Schroeder. *An Introduction to Quantum Field Theory*. Advanced Book Program, 1995.
- [23] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. “Experimental Test of Parity Conservation in Beta Decay”. *Phys. Rev.* 105 (1957), pp. 1413–1414.
- [24] E. Fermi. “An attempt of a theory of beta radiation. 1.” *Z. Phys.* 88 (1934), pp. 161–177.
- [25] K. A. Olive et al. “Review of Particle Physics”. *Chin. Phys.* C38 (2014), p. 090001.
- [26] UA1 Collaboration. “Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at $\sqrt{s} = 540$ GeV”. *Phys. Lett.* B122 (1983), pp. 103–116.
- [27] UA2 Collaboration. “Observation of Single Isolated Electrons of High Transverse Momentum in Events with Missing Transverse Energy at the CERN $\bar{p}p$ Collider”. *Phys. Lett.* B122 (1983), pp. 476–485.
- [28] UA1 Collaboration. “Experimental Observation of Lepton Pairs of Invariant Mass Around 95 GeV/ c^2 at the CERN SPS Collider”. *Phys. Lett.* B126 (1983), pp. 398–410.
- [29] UA2 Collaboration. “Evidence for $Z^0 \rightarrow e^+e^-$ at the CERN pp collider”. *Phys. Lett.* B129.1 (1983), pp. 130–140.
- [30] R. K. Ellis, W. J. Stirling, and B. R. Webber. *QCD and Collider Physics*. Cambridge University Press, 1996.
- [31] S. Dittmaier and M. Huber. “Radiative corrections to the neutral-current Drell-Yan process in the Standard Model and its minimal supersymmetric extension”. *JHEP* 01 (2010), p. 060. arXiv: 0911.2329 [hep-ph].

- [32] K. G. Wilson and J. B. Kogut. “The Renormalization group and the epsilon expansion”. *Phys. Rept.* 12 (1974), pp. 75–200.
- [33] W. Pauli and F. Villars. “On the Invariant regularization in relativistic quantum theory”. *Rev. Mod. Phys.* 21 (1949), pp. 434–444.
- [34] C. G. Bollini and J. J. Giambiagi. “Dimensional Renormalization: The Number of Dimensions as a Regularizing Parameter”. *Nuovo Cim.* B12 (1972), pp. 20–26.
- [35] G. 't Hooft and M. J. G. Veltman. “Regularization and Renormalization of Gauge Fields”. *Nucl. Phys.* B44 (1972), pp. 189–213.
- [36] O. V. Tarasov, A. A. Vladimirov, and A. Yu. Zharkov. “The Gell-Mann-Low Function of QCD in the Three Loop Approximation”. *Phys. Lett.* B93 (1980), pp. 429–432.
- [37] S. A. Larin and J. A. M. Vermaseren. “The Three loop QCD Beta function and anomalous dimensions”. *Phys. Lett.* B303 (1993), pp. 334–336. arXiv: hep-ph/9302208 [hep-ph].
- [38] T. van Ritbergen, J. A. M. Vermaseren, and S. A. Larin. “The Four loop beta function in quantum chromodynamics”. *Phys. Lett.* B400 (1997), pp. 379–384. arXiv: hep-ph/9701390 [hep-ph].
- [39] M. Czakon. “The Four-loop QCD beta-function and anomalous dimensions”. *Nucl. Phys.* B710 (2005), pp. 485–498. arXiv: hep-ph/0411261 [hep-ph].
- [40] P. A. Baikov, K. G. Chetyrkin, and J. H. Kühn. “Five-Loop Running of the QCD coupling constant” (2016). arXiv: 1606.08659 [hep-ph].
- [41] F. A. Berends and W. Giele. “The Six Gluon Process as an Example of Weyl-Van Der Waerden Spinor Calculus”. *Nucl. Phys.* B294 (1987), pp. 700–732.
- [42] M. L. Mangano, S. J. Parke, and Z. Xu. “Duality and Multi - Gluon Scattering”. *Nucl. Phys.* B298 (1988), pp. 653–672.
- [43] M. L. Mangano and S. J. Parke. “Quark - Gluon Amplitudes in the Dual Expansion”. *Nucl. Phys.* B299 (1988), pp. 673–692.

- [44] M. L. Mangano. “The Color Structure of Gluon Emission”. *Nucl. Phys.* B309 (1988), pp. 461–475.
- [45] F. A. Berends and W. T. Giele. “Recursive Calculations for Processes with n Gluons”. *Nucl. Phys.* B306 (1988), pp. 759–808.
- [46] D. A. Kosower. “Color Factorization for Fermionic Amplitudes”. *Nucl. Phys.* B315 (1989), pp. 391–418.
- [47] D. Kosower, B.-H. Lee, and V. P. Nair. “Multi Gluon Scattering: a String Based Calculation”. *Phys. Lett.* B201 (1988), pp. 85–89.
- [48] L. J. Dixon. “Calculating scattering amplitudes efficiently”. *QCD and beyond. Proceedings, Theoretical Advanced Study Institute in Elementary Particle Physics, TASI-95, Boulder, USA, June 4-30, 1995*. 1996. arXiv: hep-ph/9601359 [hep-ph].
- [49] S. Catani and M. H. Seymour. “A General algorithm for calculating jet cross-sections in NLO QCD”. *Nucl. Phys.* B485 (1997). [Erratum: *Nucl. Phys.* B510,503(1998)], pp. 291–419. arXiv: hep-ph/9605323 [hep-ph].
- [50] Z. Kunszt, A. Signer, and Z. Trocsanyi. “Singular terms of helicity amplitudes at one loop in QCD and the soft limit of the cross-sections of multiparton processes”. *Nucl. Phys.* B420 (1994), pp. 550–564. arXiv: hep-ph/9401294 [hep-ph].
- [51] W. T. Giele and E. W. N. Glover. “Higher-order corrections to jet cross sections in e^+e^- annihilation”. *Phys. Rev. D* 46 (5 1992), pp. 1980–2010.
- [52] Z. Bern and D. A. Kosower. “The Computation of loop amplitudes in gauge theories”. *Nucl. Phys.* B379 (1992), pp. 451–561.
- [53] Z. Kunszt, A. Signer, and Z. Trocsanyi. “One loop helicity amplitudes for all $2 \rightarrow 2$ processes in QCD and $N=1$ supersymmetric Yang-Mills theory”. *Nucl. Phys.* B411 (1994), pp. 397–442. arXiv: hep-ph/9305239 [hep-ph].
- [54] S. Catani, M. H. Seymour, and Z. Trocsanyi. “Regularization scheme independence and unitarity in QCD cross-sections”. *Phys. Rev.* D55 (1997), pp. 6819–6829. arXiv: hep-ph/9610553 [hep-ph].

- [55] S. Catani. “The Singular behavior of QCD amplitudes at two loop order”. *Phys. Lett.* B427 (1998), pp. 161–171. arXiv: hep-ph/9802439 [hep-ph].
- [56] G. Altarelli and G. Parisi. “Asymptotic Freedom in Parton Language”. *Nucl. Phys.* B126 (1977), pp. 298–318.
- [57] S. Catani and M. Grazzini. “Infrared factorization of tree level QCD amplitudes at the next-to-next-to-leading order and beyond”. *Nucl. Phys.* B570 (2000), pp. 287–325. arXiv: hep-ph/9908523 [hep-ph].
- [58] A. Daleo, T. Gehrmann, and D. Maître. “Antenna subtraction with hadronic initial states”. *JHEP* 04 (2007), p. 016. arXiv: hep-ph/0612257 [hep-ph].
- [59] Z. Bern, L. J. Dixon, D. C. Dunbar, and D. A. Kosower. “One loop n point gauge theory amplitudes, unitarity and collinear limits”. *Nucl. Phys.* B425 (1994), pp. 217–260. arXiv: hep-ph/9403226 [hep-ph].
- [60] Z. Bern, V. Del Duca, W. B. Kilgore, and C. R. Schmidt. “The infrared behavior of one loop QCD amplitudes at next-to-next-to leading order”. *Phys. Rev.* D60 (1999), p. 116001. arXiv: hep-ph/9903516 [hep-ph].
- [61] F. A. Berends and W. T. Giele. “Multiple Soft Gluon Radiation in Parton Processes”. *Nucl. Phys.* B313 (1989), pp. 595–633.
- [62] J. M. Campbell and E. W. N. Glover. “Double unresolved approximations to multiparton scattering amplitudes”. *Nucl. Phys.* B527 (1998), pp. 264–288. arXiv: hep-ph/9710255 [hep-ph].
- [63] R. P. Feynman. “Very high-energy collisions of hadrons”. *Phys. Rev. Lett.* 23 (1969), pp. 1415–1417.
- [64] NNPDF Collaboration. “Parton distributions with LHC data”. *Nucl. Phys.* B867 (2013), pp. 244–289. arXiv: 1207.1303 [hep-ph].
- [65] J. C. Collins, D. E. Soper, and G. F. Sterman. “Factorization of Hard Processes in QCD”. *Adv. Ser. Direct. High Energy Phys.* 5 (1989), pp. 1–91. arXiv: hep-ph/0409313 [hep-ph].
- [66] J. C. Collins, D. E. Soper, and G. F. Sterman. “Factorization for Short Distance Hadron - Hadron Scattering”. *Nucl. Phys.* B261 (1985), pp. 104–142.

- [67] A. Gehrmann-De Ridder, E. W. N. Glover, and J. Pires. “Real-Virtual corrections for gluon scattering at NNLO”. *JHEP* 02 (2012), p. 141. arXiv: 1112.3613 [hep-ph].
- [68] J. Currie, E. W. N. Glover, and S. Wells. “Infrared Structure at NNLO Using Antenna Subtraction”. *JHEP* 04 (2013), p. 066. arXiv: 1301.4693 [hep-ph].
- [69] A. Vogt, S. Moch, and J. A. M. Vermaseren. “The Three-loop splitting functions in QCD: The Singlet case”. *Nucl. Phys.* B691 (2004), pp. 129–181. arXiv: hep-ph/0404111 [hep-ph].
- [70] Y. L. Dokshitzer. “Calculation of the Structure Functions for Deep Inelastic Scattering and e^+e^- Annihilation by Perturbation Theory in Quantum Chromodynamics.” *Sov. Phys. JETP* 46 (1977). [*Zh. Eksp. Teor. Fiz.* 73,1216(1977)], pp. 641–653.
- [71] V. N. Gribov and L. N. Lipatov. “ e^+e^- pair annihilation and deep inelastic ep scattering in perturbation theory”. *Sov. J. Nucl. Phys.* 15 (1972). [*Yad. Fiz.* 15,1218(1972)], pp. 675–684.
- [72] H. Baer, J. Ohnemus, and J. F. Owens. “Next-to-leading-logarithm calculation of jet photoproduction”. *Phys. Rev. D* 40 (9 1989), pp. 2844–2855.
- [73] W. T. Giele and E. W. N. Glover. “Higher-order corrections to jet cross sections in e^+e^- annihilation”. *Phys. Rev. D* 46 (5 1992), pp. 1980–2010.
- [74] W. T. Giele, E. W. N. Glover, and D. A. Kosower. “Higher order corrections to jet cross-sections in hadron colliders”. *Nucl. Phys.* B403 (1993), pp. 633–670. arXiv: hep-ph/9302225 [hep-ph].
- [75] S. Catani and M. Grazzini. “An NNLO subtraction formalism in hadron collisions and its application to Higgs boson production at the LHC”. *Phys. Rev. Lett.* 98 (2007), p. 222002. arXiv: hep-ph/0703012 [hep-ph].
- [76] R. Boughezal, C. Focke, X. Liu, and F. Petriello. “ W -boson production in association with a jet at next-to-next-to-leading order in perturbative QCD”. *Phys. Rev. Lett.* 115.6 (2015), p. 062002. arXiv: 1504.02131 [hep-ph].

- [77] C. W. Bauer, S. Fleming, D. Pirjol, and I. W. Stewart. “An Effective field theory for collinear and soft gluons: Heavy to light decays”. *Phys. Rev. D* 63 (2001), p. 114020. arXiv: hep-ph/0011336 [hep-ph].
- [78] E. C. Poggio, H. R. Quinn, and S. Weinberg. “Smearing method in the quark model”. *Phys. Rev. D* 13 (7 1976), pp. 1958–1968.
- [79] Y. I. Azimov, Y. L. Dokshitzer, V. A. Khoze, and S. I. Troyan. “Similarity of Parton and Hadron Spectra in QCD Jets”. *Z. Phys. C* 27 (1985), pp. 65–72.
- [80] J. E. Huth et al. “Toward a standardization of jet definitions”. *1990 DPF Summer Study on High-energy Physics: Research Directions for the Decade (Snowmass 90) Snowmass, Colorado, June 25-July 13, 1990*. 1990.
- [81] G. C. Blazey et al. “Run II jet physics”. *QCD and weak boson physics in Run II. Proceedings, Batavia, USA, March 4-6, June 3-4, November 4-6, 1999*. 2000, pp. 47–77. arXiv: hep-ex/0005012 [hep-ex].
- [82] G. P. Salam and G. Soyez. “A Practical Seedless Infrared-Safe Cone jet algorithm”. *JHEP* 05 (2007), p. 086. arXiv: 0704.0292 [hep-ph].
- [83] S. Catani, Y. L. Dokshitzer, M. H. Seymour, and B. R. Webber. “Longitudinally invariant K_t clustering algorithms for hadron hadron collisions”. *Nucl. Phys. B* 406 (1993), pp. 187–224.
- [84] S. D. Ellis and D. E. Soper. “Successive combination jet algorithm for hadron collisions”. *Phys. Rev. D* 48 (1993), pp. 3160–3166. arXiv: hep-ph/9305266 [hep-ph].
- [85] Y. L. Dokshitzer, G. D. Leder, S. Moretti, and B. R. Webber. “Better jet clustering algorithms”. *JHEP* 08 (1997), p. 001. arXiv: hep-ph/9707323 [hep-ph].
- [86] M. Wobisch and T. Wengler. “Hadronization corrections to jet cross-sections in deep inelastic scattering”. *Monte Carlo generators for HERA physics. Proceedings, Workshop, Hamburg, Germany, 1998-1999*. 1998. arXiv: hep-ph/9907280 [hep-ph].
- [87] M. Cacciari, G. P. Salam, and G. Soyez. “The Anti- k_T jet clustering algorithm”. *JHEP* 04 (2008), p. 063. arXiv: 0802.1189 [hep-ph].

- [88] OPAL Collaboration. “QCD studies using a cone based jet finding algorithm for e^+e^- collisions at LEP”. *Z. Phys.* C63 (1994), pp. 197–212.
- [89] M. H. Seymour. “Jet shapes in hadron collisions: Higher orders, resummation and hadronization”. *Nucl. Phys.* B513 (1998), pp. 269–300. arXiv: hep-ph/9707338 [hep-ph].
- [90] D. A. Kosower. “Antenna factorization of gauge theory amplitudes”. *Phys. Rev.* D57 (1998), pp. 5410–5416. arXiv: hep-ph/9710213 [hep-ph].
- [91] D. A. Kosower. “Multiple singular emission in gauge theories”. *Phys. Rev.* D67 (2003), p. 116003. arXiv: hep-ph/0212097 [hep-ph].
- [92] E. W. N. Glover and J. Pires. “Antenna subtraction for gluon scattering at NNLO”. *JHEP* 06 (2010), p. 096. arXiv: 1003.2824 [hep-ph].
- [93] A. Gehrmann-De Ridder, T. Gehrmann, and G. Heinrich. “Four particle phase space integrals in massless QCD”. *Nucl. Phys.* B682 (2004), pp. 265–288. arXiv: hep-ph/0311276 [hep-ph].
- [94] A. Gehrmann-De Ridder, T. Gehrmann, and M. Ritzmann. “Antenna subtraction at NNLO with hadronic initial states: double real initial-initial configurations”. *JHEP* 10 (2012), p. 047. arXiv: 1207.5779 [hep-ph].
- [95] A. Daleo, A. Gehrmann-De Ridder, T. Gehrmann, and G. Luisoni. “Antenna subtraction at NNLO with hadronic initial states: initial-final configurations”. *JHEP* 01 (2010), p. 118. arXiv: 0912.0374 [hep-ph].
- [96] R. Boughezal, A. Gehrmann-De Ridder, and M. Ritzmann. “Antenna subtraction at NNLO with hadronic initial states: double real radiation for initial-initial configurations with two quark flavours”. *JHEP* 02 (2011), p. 098. arXiv: 1011.6631 [hep-ph].
- [97] T. Gehrmann and P. F. Monni. “Antenna subtraction at NNLO with hadronic initial states: real-virtual initial-initial configurations”. *JHEP* 12 (2011), p. 049. arXiv: 1107.4037 [hep-ph].
- [98] T. Gleisberg and F. Krauss. “Automating dipole subtraction for QCD NLO calculations”. *Eur. Phys. J.* C53 (2008), pp. 501–523. arXiv: 0709.2881 [hep-ph].

- [99] R. Frederix, S. Frixione, F. Maltoni, and T. Stelzer. “Automation of next-to-leading order computations in QCD: The FKS subtraction”. *JHEP* 10 (2009), p. 003. arXiv: 0908.4272 [hep-ph].
- [100] A. Gehrmann-De Ridder, T. Gehrmann, E. W. N. Glover, and G. Heinrich. “Infrared structure of $e^+e^- \rightarrow 3$ jets at NNLO”. *JHEP* 11 (2007), p. 058. arXiv: 0710.0346 [hep-ph].
- [101] J. Pires and E. W. N. Glover. “Double real radiation corrections to gluon scattering at NNLO”. *Nucl. Phys. Proc. Suppl.* 205-206 (2010), pp. 176–181. arXiv: 1006.1849 [hep-ph].
- [102] Z. Bern, L. J. Dixon, and D. A. Kosower. “One loop amplitudes for e^+e^- to four partons”. *Nucl. Phys.* B513 (1998), pp. 3–86. arXiv: hep-ph/9708239 [hep-ph].
- [103] S. D. Drell and T.-M. Yan. “Massive Lepton-Pair Production in Hadron-Hadron Collisions at High Energies”. *Phys. Rev. Lett.* 25 (5 1970), pp. 316–320.
- [104] J. H. Christenson, G. S. Hicks, L. M. Lederman, P. J. Limon, B. G. Pope, and E. Zavattini. “Observation of Muon Pairs in High-Energy Hadron Collisions”. *Phys. Rev. D* 8 (7 1973), pp. 2016–2034.
- [105] J. M. Campbell, J. W. Huston, and W. J. Stirling. “Hard Interactions of Quarks and Gluons: A Primer for LHC Physics”. *Rept. Prog. Phys.* 70 (2007), p. 89. arXiv: hep-ph/0611148 [hep-ph].
- [106] H. Baer et al. “Low-energy supersymmetry phenomenology” (1995). arXiv: hep-ph/9503479 [hep-ph].
- [107] B. C. Allanach, C. G. Lester, M. A. Parker, and B. R. Webber. “Measuring sparticle masses in nonuniversal string inspired models at the LHC”. *JHEP* 09 (2000), p. 004. arXiv: hep-ph/0007009 [hep-ph].
- [108] G. Altarelli, R. K. Ellis, and G. Martinelli. “Leptoproduction and Drell-Yan Processes Beyond the Leading Approximation in Chromodynamics”. *Nucl. Phys.* B143 (1978). [Erratum: *Nucl. Phys.*B146,544(1978)], p. 521.
- [109] J. Kubar-Andre and F. E. Paige. “Gluon corrections to the Drell-Yan model”. *Phys. Rev. D* 19 (1 1979), pp. 221–229.

- [110] K. Harada, T. Kaneko, and N. Sakai. “Hadronic Lepton Pair Production Beyond the Leading Order in Perturbative QCD”. *Nucl. Phys.* B155 (1979). [Erratum: *Nucl. Phys.*B165,545(1980)], p. 169.
- [111] B. Humpert and W. L. Van Neerven. “Ambiguities in the infrared regularization of QCD”. *Phys. Lett.* B84 (1979). [Erratum: *Phys. Lett.*B85,471(1979)], p. 327.
- [112] B. Humpert and W. L. van Neerven. “How to Regularize the Infrared and Mass Singularities in QCD”. *Phys. Lett.* B89 (1979), pp. 69–75.
- [113] B. Humpert and W. L. van Neerven. “Infrared and Mass Regularization in Affine Field Theories 2. QCD”. *Nucl. Phys.* B184 (1981), pp. 225–268.
- [114] J. Kubar, M. Le Bellac, J. L. Meunier, and G. Plaut. “QCD Corrections to the Drell-Yan Mechanism and the Pion Structure Function”. *Nucl. Phys.* B175 (1980), pp. 251–275.
- [115] R. Hamberg, W. L. van Neerven, and T. Matsuura. “A Complete calculation of the order α_s^2 correction to the Drell-Yan K factor”. *Nucl. Phys.* B359 (1991). [Erratum: *Nucl. Phys.*B644,403(2002)], pp. 343–405.
- [116] W. L. van Neerven and E. B. Zijlstra. “The $\mathcal{O}(\alpha_s^2)$ corrected Drell-Yan K factor in the DIS and $\overline{\text{MS}}$ scheme”. *Nucl. Phys.* B382 (1992). [Erratum: *Nucl. Phys.*B680, 513(2004)], pp. 11–62.
- [117] U. Baur, O. Brein, W. Hollik, C. Schappacher, and D. Wackerroth. “Electroweak radiative corrections to neutral current Drell-Yan processes at hadron colliders”. *Phys. Rev.* D65 (2002), p. 033007. arXiv: hep-ph/0108274 [hep-ph].
- [118] L. Barze, G. Montagna, P. Nason, O. Nicrosini, F. Piccinini, and A. Vicini. “Neutral current Drell-Yan with combined QCD and electroweak corrections in the POWHEG BOX”. *Eur. Phys. J.* C73.6 (2013), p. 2474. arXiv: 1302.4606 [hep-ph].
- [119] G. Bozzi, S. Catani, G. Ferrera, D. de Florian, and M. Grazzini. “Transverse-momentum resummation: A Perturbative study of Z production at the Tevatron”. *Nucl. Phys.* B815 (2009), pp. 174–197. arXiv: 0812.2862 [hep-ph].

- [120] G. Bozzi, S. Catani, G. Ferrera, D. de Florian, and M. Grazzini. “Production of Drell-Yan lepton pairs in hadron collisions: Transverse-momentum resummation at next-to-next-to-leading logarithmic accuracy”. *Phys. Lett.* B696 (2011), pp. 207–213. arXiv: 1007.2351 [hep-ph].
- [121] S. Dittmaier, A. Huss, and C. Schwinn. “Mixed QCD-electroweak $\mathcal{O}(\alpha_s\alpha)$ corrections to Drell-Yan processes in the resonance region: pole approximation and non-factorizable corrections”. *Nucl. Phys.* B885 (2014), pp. 318–372. arXiv: 1403.3216 [hep-ph].
- [122] S. Dittmaier, A. Huss, and C. Schwinn. “Dominant mixed QCD-electroweak $\mathcal{O}(\alpha_s\alpha)$ corrections to Drell-Yan processes in the resonance region”. *Nucl. Phys.* B904 (2016), pp. 216–252. arXiv: 1511.08016 [hep-ph].
- [123] W. T. Giele, E. W. N. Glover, and D. A. Kosower. “Higher order corrections to jet cross-sections in hadron colliders”. *Nucl. Phys.* B403 (1993), pp. 633–670. arXiv: hep-ph/9302225 [hep-ph].
- [124] J. M. Campbell and R. K. Ellis. “Next-to-leading order corrections to $W + 2$ jet and $Z + 2$ jet production at hadron colliders”. *Phys. Rev.* D65 (2002), p. 113007. arXiv: hep-ph/0202176 [hep-ph].
- [125] J. M. Campbell, R. K. Ellis, and D. L. Rainwater. “Next-to-leading order QCD predictions for $W + 2$ jet and $Z + 2$ jet production at the CERN LHC”. *Phys. Rev.* D68 (2003), p. 094021. arXiv: hep-ph/0308195 [hep-ph].
- [126] C. F. Berger, Z. Bern, L. J. Dixon, F. Febres Cordero, D. Forde, T. Gleisberg, H. Ita, D. A. Kosower, and D. Maitre. “Next-to-Leading Order QCD Predictions for $Z, \gamma^* + 3$ -Jet Distributions at the Tevatron”. *Phys. Rev.* D82 (2010), p. 074002. arXiv: 1004.1659 [hep-ph].
- [127] H. Ita, Z. Bern, L. J. Dixon, F. Febres Cordero, D. A. Kosower, and D. Maitre. “Precise Predictions for $Z + 4$ Jets at Hadron Colliders”. *Phys. Rev.* D85 (2012), p. 031501. arXiv: 1108.2229 [hep-ph].
- [128] J. H. Kuhn, A. Kulesza, S. Pozzorini, and M. Schulze. “One-loop weak corrections to hadronic production of Z bosons at large transverse momenta”. *Nucl. Phys.* B727 (2005), pp. 368–394. arXiv: hep-ph/0507178 [hep-ph].

- [129] A. Denner, S. Dittmaier, T. Kasprzik, and A. Muck. “Electroweak corrections to dilepton + jet production at hadron colliders”. *JHEP* 06 (2011), p. 069. arXiv: 1103.0914 [hep-ph].
- [130] S. Kallweit, J. M. Lindert, P. Maierhöfer, S. Pozzorini, and M. Schönherr. “NLO QCD+EW predictions for $V + \text{jets}$ including off-shell vector-boson decays and multijet merging”. *JHEP* 04 (2016), p. 021. arXiv: 1511.08692 [hep-ph].
- [131] S. Moch, P. Uwer, and S. Weinzierl. “Two loop amplitudes with nested sums: Fermionic contributions to $e^+e^- \rightarrow q\bar{q}g$ ”. *Phys. Rev. D* 66 (2002), p. 114001. arXiv: hep-ph/0207043 [hep-ph].
- [132] L. W. Garland, T. Gehrmann, E. W. N. Glover, A. Koukoutsakis, and E. Remiddi. “The Two loop QCD matrix element for $e^+e^- \rightarrow 3 \text{ jets}$ ”. *Nucl. Phys. B* 627 (2002), pp. 107–188. arXiv: hep-ph/0112081 [hep-ph].
- [133] L. W. Garland, T. Gehrmann, E. W. N. Glover, A. Koukoutsakis, and E. Remiddi. “Two loop QCD helicity amplitudes for $e^+e^- \rightarrow \text{three jets}$ ”. *Nucl. Phys. B* 642 (2002), pp. 227–262. arXiv: hep-ph/0206067 [hep-ph].
- [134] T. Gehrmann and L. Tancredi. “Two-loop QCD helicity amplitudes for $q\bar{q} \rightarrow W^\pm\gamma$ and $q\bar{q} \rightarrow Z^0\gamma$ ”. *JHEP* 02 (2012), p. 004. arXiv: 1112.1531 [hep-ph].
- [135] E. W. N. Glover and D. J. Miller. “The One loop QCD corrections for $\gamma^* \rightarrow Q\bar{Q}q\bar{q}$ ”. *Phys. Lett. B* 396 (1997), pp. 257–263. arXiv: hep-ph/9609474 [hep-ph].
- [136] Z. Bern, L. J. Dixon, D. A. Kosower, and S. Weinzierl. “One loop amplitudes for $e^+e^- \rightarrow \bar{q}q\bar{Q}Q$ ”. *Nucl. Phys. B* 489 (1997), pp. 3–23. arXiv: hep-ph/9610370 [hep-ph].
- [137] J. M. Campbell, E. W. N. Glover, and D. J. Miller. “The One loop QCD corrections for $\gamma^* \rightarrow q\bar{q}gg$ ”. *Phys. Lett. B* 409 (1997), pp. 503–508. arXiv: hep-ph/9706297 [hep-ph].
- [138] K. Hagiwara and D. Zeppenfeld. “Amplitudes for Multiparton Processes Involving a Current at e^+e^- , $e^\pm p$, and Hadron Colliders”. *Nucl. Phys. B* 313 (1989), pp. 560–594.

- [139] F. A. Berends, W. T. Giele, and H. Kuijf. “Exact Expressions for Processes Involving a Vector Boson and Up to Five Partons”. *Nucl. Phys.* B321 (1989), pp. 39–82.
- [140] N. K. Falck, D. Graudenz, and G. Kramer. “Cross-section for Five Jet Production in e^+e^- Annihilation”. *Nucl. Phys.* B328 (1989), pp. 317–341.
- [141] R. Boughezal, J. M. Campbell, R. K. Ellis, C. Focke, W. T. Giele, X. Liu, and F. Petriello. “Z-boson production in association with a jet at next-to-next-to-leading order in perturbative QCD”. *Phys. Rev. Lett.* 116.15 (2016), p. 152001. arXiv: 1512.01291 [hep-ph].
- [142] ATLAS Collaboration. “Measurement of the Z/γ^* boson transverse momentum distribution in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector”. *JHEP* 09 (2014), p. 145. arXiv: 1406.3660 [hep-ex].
- [143] ATLAS Collaboration. “Measurement of the transverse momentum and ϕ_η^* distributions of Drell-Yan lepton pairs in proton-proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector” (2015). arXiv: 1512.02192 [hep-ex].
- [144] CMS Collaboration. “Measurement of the Rapidity and Transverse Momentum Distributions of Z Bosons in pp Collisions at $\sqrt{s} = 7$ TeV”. *Phys. Rev.* D85 (2012), p. 032002. arXiv: 1110.4973 [hep-ex].
- [145] CMS Collaboration. “Measurement of the Z boson differential cross section in transverse momentum and rapidity in proton-proton collisions at 8 TeV”. *Phys. Lett.* B749 (2015), pp. 187–209. arXiv: 1504.03511 [hep-ex].
- [146] LHCb Collaboration. “Measurement of forward W and Z boson production in pp collisions at $\sqrt{s} = 8$ TeV”. *JHEP* 01 (2016), p. 155. arXiv: 1511.08039 [hep-ex].
- [147] J. A. M. Vermaseren. “New features of FORM” (2000). arXiv: math-ph/0010025 [math-ph].
- [148] R. Kleiss, W. J. Stirling, and S. D. Ellis. “A New Monte Carlo Treatment of Multiparticle Phase Space at High-energies”. *Comput. Phys. Commun.* 40 (1986), p. 359.

- [149] G. P. Lepage. “A New Algorithm for Adaptive Multidimensional Integration”. *J. Comput. Phys.* 27 (1978), p. 192.
- [150] P. J. W. Faulkner et al. “GridPP: Development of the UK computing Grid for particle physics”. *J. Phys.* G32 (2006), N1–N20.
- [151] R. Kleiss and R. Pittau. “Weight optimization in multichannel Monte Carlo”. *Comput. Phys. Commun.* 83 (1994), pp. 141–146. arXiv: hep-ph/9405257 [hep-ph].
- [152] J. J. van der Bij and E. W. N. Glover. “Z Boson Production and Decay via Gluons”. *Nucl. Phys.* B313 (1989), pp. 237–257.
- [153] T. Carli, D. Clements, A. Cooper-Sarkar, C. Gwenlan, G. P. Salam, F. Siegert, P. Starovoitov, and M. Sutton. “A posteriori inclusion of parton density functions in NLO QCD final-state calculations at hadron colliders: The APPLGRID Project”. *Eur. Phys. J.* C66 (2010), pp. 503–524. arXiv: 0911.2985 [hep-ph].
- [154] D. Britzger, K. Rabbertz, F. Stober, and M. Wobisch. “New features in version 2 of the fastNLO project”. *Proceedings, 20th International Workshop on Deep-Inelastic Scattering and Related Subjects (DIS 2012)*. 2012, pp. 217–221. arXiv: 1208.3641 [hep-ph].
- [155] Z. Bern, L. J. Dixon, F. Febres Cordero, S. Höche, H. Ita, D. A. Kosower, and D. Maître. “Ntuples for NLO Events at Hadron Colliders”. *Comput. Phys. Commun.* 185 (2014), pp. 1443–1460. arXiv: 1310.7439 [hep-ph].
- [156] S. Catani and B. R. Webber. “Infrared safe but infinite: Soft gluon divergences inside the physical region”. *JHEP* 10 (1997), p. 005. arXiv: hep-ph/9710333 [hep-ph].
- [157] NNPDF Collaboration. “Parton distributions for the LHC Run II”. *JHEP* 04 (2015), p. 040. arXiv: 1410.8849 [hep-ph].
- [158] J. C. Collins, D. E. Soper, and G. F. Sterman. “Transverse Momentum Distribution in Drell-Yan Pair and W and Z Boson Production”. *Nucl. Phys.* B250 (1985), pp. 199–224.

-
- [159] T. Becher and M. Neubert. “Drell-Yan Production at Small q_T , Transverse Parton Distributions and the Collinear Anomaly”. *Eur. Phys. J. C* 71 (2011), p. 1665. arXiv: 1007.4005 [hep-ph].
- [160] X. Chen, T. Gehrmann, E. W. N. Glover, and M. Jaquier. “Higgs plus one jet production at NNLO”. *Proceedings, 12th International Symposium on Radiative Corrections (Radcor 2015) and LoopFest XIV (Radiative Corrections for the LHC and Future Colliders)*. 2016. arXiv: 1604.04085 [hep-ph].
- [161] J. R. Andersen et al. “Les Houches 2015: Physics at TeV Colliders Standard Model Working Group Report”. *9th Les Houches Workshop on Physics at TeV Colliders (PhysTeV 2015) Les Houches, France, June 1-19, 2015*. 2016. arXiv: 1605.04692 [hep-ph].
- [162] G. Heinrich, M. Johnson, and D. Maître. “N(N)LO event files: applications and prospects”. *13th DESY Workshop on Elementary Particle Physics: Loops and Legs in Quantum Field Theory (LL2016) Leipzig, Germany, April 24-29, 2016*. 2016. arXiv: 1607.06259 [hep-ph].
- [163] R. Boughezal, J. M. Campbell, R. K. Ellis, C. Focke, W. Giele, X. Liu, F. Petriello, and C. Williams. “Color singlet production at NNLO in MCFM” (2016). arXiv: 1605.08011 [hep-ph].

Appendix A

X_3^0 Antenna functions

In this appendix we list the explicit X_3^0 unintegrated antenna functions given in terms of invariants. We use the naming convention defined in Tab. 2.1. Explicit formulae for all other antennae, including integrated antenna functions, can be found in Ref. [11].

A.1 $\gamma^* \rightarrow q\bar{q}$ antenna functions

At NLO only one antenna function is derived from the process $\gamma^* \rightarrow q\bar{q}$, namely [11],

$$A_3^0(1_q, 3_g, 2_{\bar{q}}) = \frac{1}{s_{123}} \left(\frac{s_{13}}{s_{23}} + \frac{s_{23}}{s_{13}} + 2 \frac{s_{12}s_{123}}{s_{13}s_{23}} \right), \quad (\text{A.1.1})$$

which contains the divergent limits,

- 3_g soft,
- $1_q || 3_g$ collinear,
- $2_{\bar{q}} || 3_g$ collinear.

From this antenna function the sub-antenna a_3^0 can be defined,

$$a_3^0(1_q, 3_g, 2_{\bar{q}}) = \frac{1}{s_{123}} \left(\frac{s_{23}}{s_{13}} + 2 \frac{s_{12}s_{123}}{(s_{13} + s_{23})s_{13}} \right), \quad (\text{A.1.2})$$

where,

$$A_3^0(1_q, 3_g, 2_{\bar{q}}) = a_3^0(1_q, 3_g, 2_{\bar{q}}) + a_3^0(2_{\bar{q}}, 3_g, 1_{\bar{q}}). \quad (\text{A.1.3})$$

The a_3^0 sub-antenna function contains the $1_q || 3_g$ collinear limit and partially the 3 soft limit. The full 3 soft limit is only recovered when two a_3^0 antenna functions are used

in conjunction with each other. As a result of this the a_3^0 antenna function is typically only useful when the gluon is in the initial state.

A.2 $\tilde{\chi} \rightarrow \tilde{g}g$ antenna functions

There are two tree level three parton antenna functions derived from the neutralino decay process,

- $\tilde{\chi} \rightarrow \tilde{g}gg$,
- $\tilde{\chi} \rightarrow \tilde{g}q\bar{q}'$,

which are denoted by $D_3^0(1_q, 3_g, 4_g)$ and $E_3^0(1_q, 3_{q'}, 4_{\bar{q}'})$ respectively. The explicit formula for $D_3^0(1_q, 3_g, 4_g)$ is given in Eq. (2.1.35) and the symmetric decomposition of the antenna into $d_3^0(1_q, 3_g, 4_g)$ sub-antennae is given in Eq. (2.1.38).

In configurations where a gluon is in the initial state, the flavour preserving $d_{3,g \rightarrow g}^0(3_q, 4_g, \hat{1}_g)$ sub-antenna is given by,

$$d_{3,g \rightarrow g}^0(3_q, 4_g, \hat{1}_g) = \frac{1}{s_{\hat{1}34}^2} \left(\frac{s_{\hat{1}4}^2}{s_{34}} + 4 \frac{s_{\hat{1}3}^2}{s_{34}} + 4 \frac{s_{\hat{1}3}^2}{s_{\hat{1}4}} + 4 \frac{s_{34}^2}{s_{\hat{1}4}} + 3 \frac{s_{\hat{1}4} s_{\hat{1}3}}{s_{34}} + 2 \frac{s_{\hat{1}3}^3}{s_{34} s_{\hat{1}4}} \right. \\ \left. + 2 \frac{s_{34}^3}{s_{\hat{1}4} (s_{\hat{1}3} + s_{\hat{1}4})} + 6 \frac{s_{34} s_{\hat{1}3}}{s_{\hat{1}4}} + 6 s_{\hat{1}4} + 9 (s_{\hat{1}3} + s_{34}) \right), \quad (\text{A.2.1})$$

which contains the divergent limits,

- 4_g soft,
- $3_q || 4_g$ collinear,
- $\hat{1}_q || 4_g$ collinear.

The flavour changing sub-antenna, $d_{3,g \rightarrow q}^0(3_q, \hat{1}_g, 4_g)$, is given by,

$$d_{3,g \rightarrow q}^0(3_q, \hat{1}_g, 4_g) = \frac{1}{s_{\hat{1}34}^2} \left(\frac{s_{\hat{1}4}^2}{s_{\hat{1}3}} + 4 \frac{s_{34}^2}{s_{\hat{1}3}} + 3 \frac{s_{34} s_{\hat{1}4}}{s_{\hat{1}3}} + 2 \frac{s_{34}^3}{s_{\hat{1}3} (s_{\hat{1}3} + s_{34})} \right), \quad (\text{A.2.2})$$

which only contains the $\hat{1}_g || 3_q$ collinear limit.

The $E_3^0(1_q, 3_{q'}, 4_{\bar{q}'})$ antenna function is given by [11],

$$E_3^0(1_q, 3_{q'}, 4_{\bar{q}'}) = \frac{1}{s_{\hat{1}34}^2} \left(\frac{s_{\hat{1}3}^2 + s_{\hat{1}4}^2}{s_{34}} + s_{13} + s_{14} \right), \quad (\text{A.2.3})$$

and contains only the $3_{q'} || 4_{\bar{q}'}$ collinear limit.

A.3 $H \rightarrow gg$ antenna functions

There are two tree level three parton antenna functions derived from the Higgs decay process,

- $H \rightarrow ggg$,
- $H \rightarrow gq\bar{q}$,

which are denoted by $F_3^0(1_g, 2_g, 3_g)$ and $G_3^0(1_g, 2_q, 3_{\bar{q}})$ respectively. The $F_3^0(1_g, 2_g, 3_g)$ antenna function is given by,

$$F_3^0(1_g, 2_g, 3_g) = \frac{2}{s_{123}^2} \left(\frac{s_{123}^2 s_{12}}{s_{13} s_{23}} + \frac{s_{123}^2 s_{13}}{s_{12} s_{23}} + \frac{s_{123}^2 s_{23}}{s_{12} s_{13}} + \frac{s_{12} s_{13}}{s_{23}} + \frac{s_{12} s_{23}}{s_{13}} + \frac{s_{13} s_{23}}{s_{12}} + 4s_{123} \right), \quad (\text{A.3.1})$$

which contains all possible soft and collinear limits between the gluons. This can be decomposed into the sub-antenna $f_3^0(1_g, 3_g, 2_g)$ [11],

$$f_3^0(1_g, 3_g, 2_g) = \frac{1}{s_{123}^2} \left(2 \frac{s_{123}^2 s_{12}}{s_{13} s_{23}} + \frac{s_{12} s_{13}}{s_{23}} + \frac{s_{12} s_{23}}{s_{13}} + \frac{8}{3} s_{123} \right), \quad (\text{A.3.2})$$

where,

$$F_3^0(1_g, 2_g, 3_g) = f_3^0(1_g, 3_g, 2_g) + f_3^0(3_g, 2_g, 1_g) + f_3^0(2_g, 1_g, 3_g). \quad (\text{A.3.3})$$

The sub-antenna $f_3^0(1_g, 3_g, 2_g)$ contains the limits,

- 3_g soft,
- partial $1_g || 3_g$ collinear limit,
- partial $2_g || 3_g$ collinear limit.

The $G_3^0(1_g, 3_q, 4_{\bar{q}})$ antenna function is given by [11],

$$G_3^0(1_g, 3_q, 4_{\bar{q}}) = \frac{1}{s_{134}^2} \left(\frac{s_{13}^2 + s_{14}^2}{s_{34}} \right), \quad (\text{A.3.4})$$

which similarly to the $E_3^0(1_q, 3_{q'}, 4_{\bar{q}'})$ antenna function only contains the $3_q || 4_{\bar{q}}$ collinear limit.

Appendix B

Validation plots and distributions

B.1 NLO validation

Two separate NLO calculations are required when considering the inclusive Z +jet NNLO computation,

- Z +jet at NLO,
- Z +2 jets at NLO,

both of which can be validated against independent calculations. These NLO calculations can be found within the code MCFM [163], where we can compare both the total cross section and differential distributions for a wide range of observables.

Unless explicitly stated otherwise, the validation plots and distributions generated in sections B.1.1 and B.2 were generated using the setup described in section 6.1 with a scale choice of $\mu_F = \mu_R = M_Z$. In section B.1.2 we require at least two resolved jets. We use the same setup as above, however we also include an additional requirement of a second jet with transverse momentum of at least 30 GeV ($p_T^{\text{jet}_2} > 30$ GeV).

B.1.1 Z +jet at NLO

Testing the calculation of Z +jet production at NLO against publicly available tools is an important step in performing the Z +jet calculation at NNLO. It tests the implementation of the tree level 4 parton matrix elements, and also the 3 parton one-loop

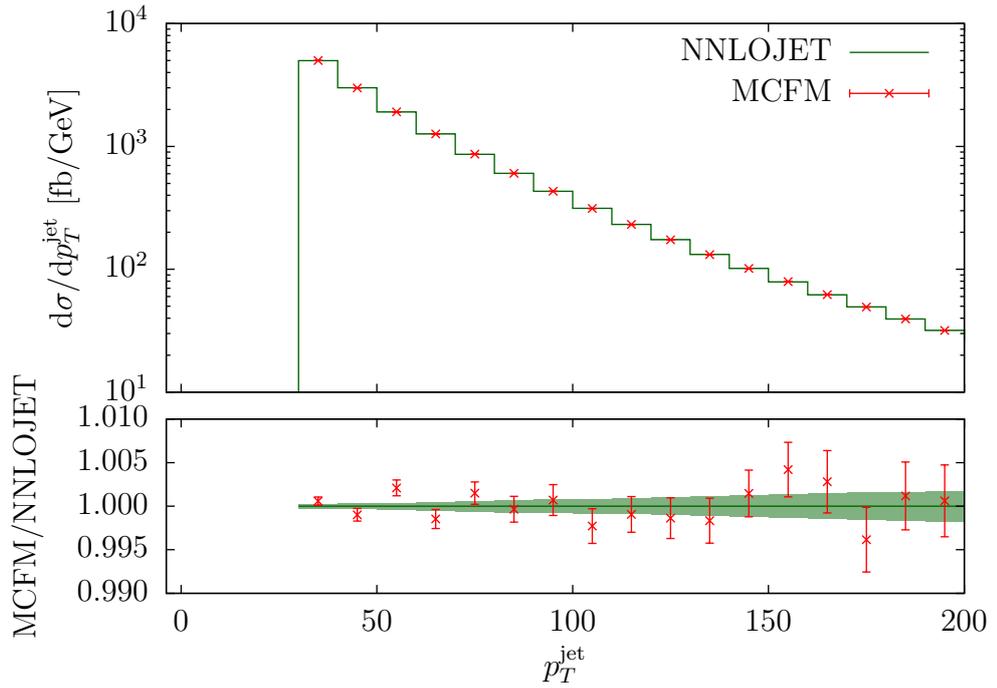


Figure B.1: The transverse momentum distribution of the leading jet in inclusive Z +jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical.

matrix elements. The volume of the 3 parton and 4 parton phase spaces are also thoroughly validated. Furthermore, a lot of the code structure of the virtual integration at NLO is similar to that required for the double virtual and real virtual contributions at NNLO. Clearly any potential problems that could arise in an NLO calculation would impact the calculation of the NNLO corrections.

The NLO calculation for Z +jet production using the NNLOJET code gives the total cross section,

$$\sigma_{\text{NLO}}^{ZJ} = (144.361 \pm 0.013) \text{ pb},$$

where the uncertainty is a statistical uncertainty from the numerical integration. Performing the same calculation using MCFM gives,

$$\sigma_{\text{NLO}}^{ZJ} = (144.372 \pm 0.023) \text{ pb}.$$

We can conclude that there is agreement up to 0.01% between the NNLOJET implementation and the MCFM implementation of the NLO calculation for the total cross section.

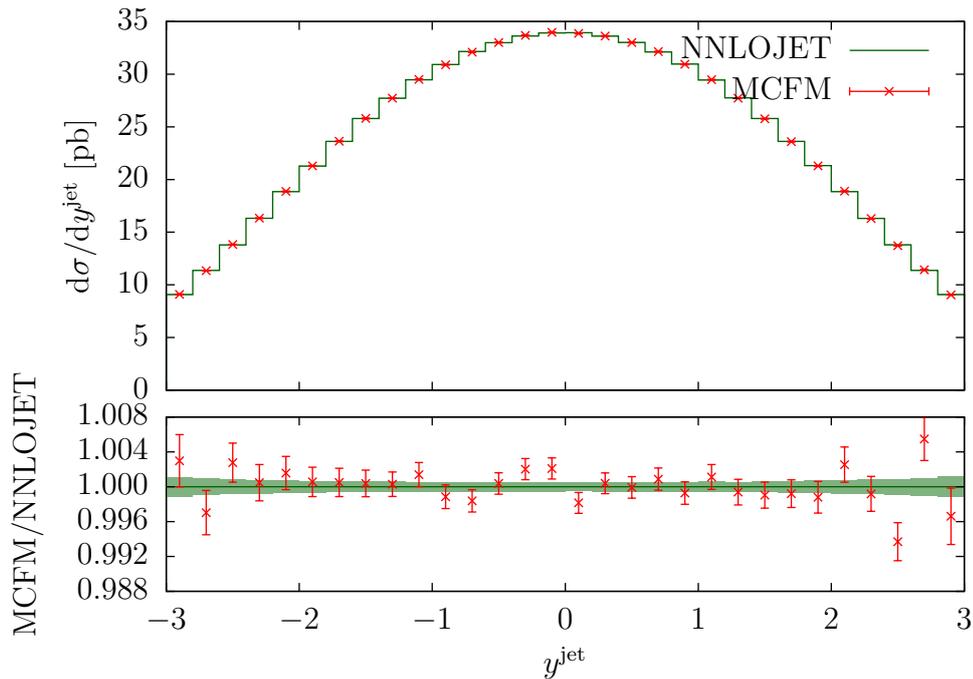


Figure B.2: The rapidity distribution of the leading jet in inclusive Z +jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical.

Clearly it is also important to validate a range of observables, particularly those observables presented in section 6.1. Fig. B.1 shows the transverse momentum of the leading jet in inclusive Z +jet production at NLO, calculated using both NNLOJET (green) and MCFM (red). Per-mille level agreement can be observed between the implementations across a wide range of transverse momentum.

Similarly Fig. B.2 shows the rapidity distribution of the leading jet, calculated using both NNLOJET and MCFM. Again we see per-mille agreement between the two codes.

Fig. B.3 shows the transverse momentum distribution of the Z boson in Z +jet production. Both codes observe the same Sudakov phenomenon [156], as discussed in section 6.1. In the bin close to the jet cut (30 GeV) there is marginal tension between the two codes. This is a numerically unstable region of the phase space where either or both calculations are underestimating the statistical uncertainties associated with this bin. Nevertheless, given the limited predictive power of the fixed order calculation in this region it has limited impact on the physical conclusions drawn from using MCFM or NNLOJET to perform the NLO calculation. Other than this particular bin, per-mille

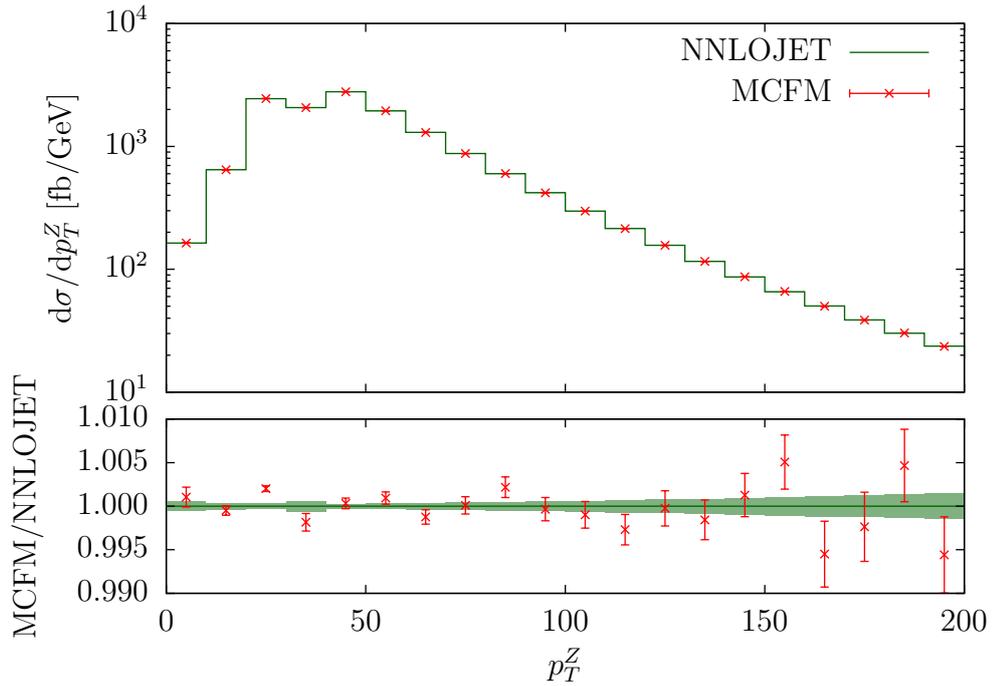


Figure B.3: The transverse momentum distribution of the Z boson in inclusive Z +jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical.

level agreement is observed across the entire range of transverse momentum.

Fig. B.4 shows the rapidity distribution of the Z boson in Z +jet production. Again, excellent agreement is observed between the two implementations. Even in the very forward regions where the cross section is small, both NNLOJET and MCFM agree to per-cent level accuracy.

Fig. B.5 shows the transverse momentum distribution for the negatively charged lepton for inclusive Z +jet production. Again per-mille level agreement is found between MCFM and NNLOJET across the entire range of transverse momentum. Similar conclusions can be drawn from the rapidity distribution of the negatively charged lepton shown in Fig. B.6.

B.1.2 Z +2 jets at NLO

As in section B.1.1, a calculation of Z +2 jets at NLO can be performed using NNLOJET and compared to a publicly available implementation of the calculation found in MCFM. This tests the implementation of the tree level 5 parton matrix elements and also the 4 parton one-loop matrix elements. The volume of the 4 parton and 5 parton

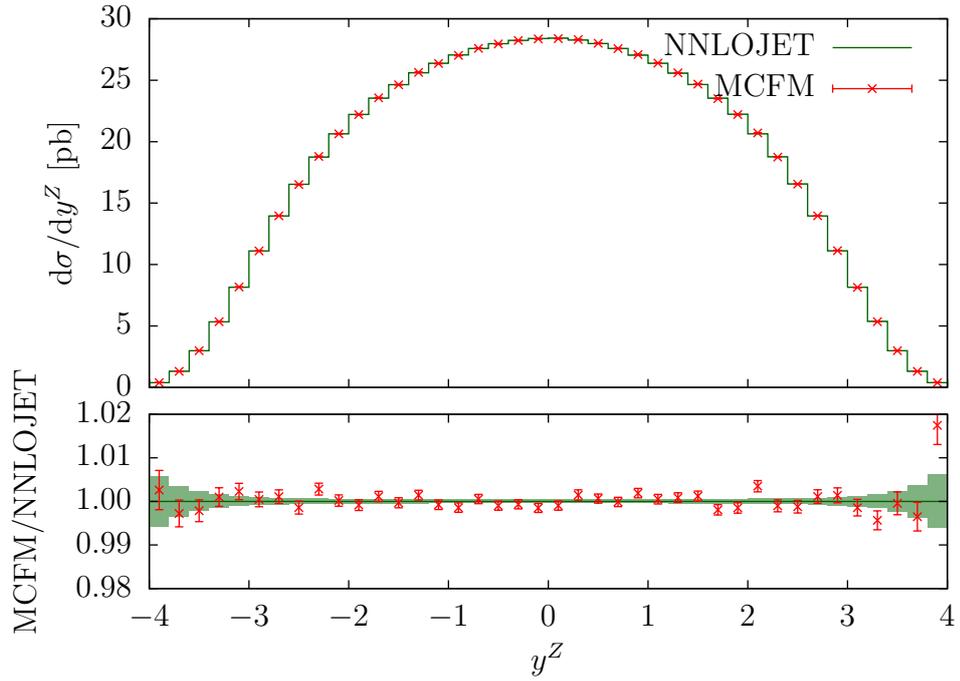


Figure B.4: The rapidity distribution of the Z boson in inclusive Z +jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical.

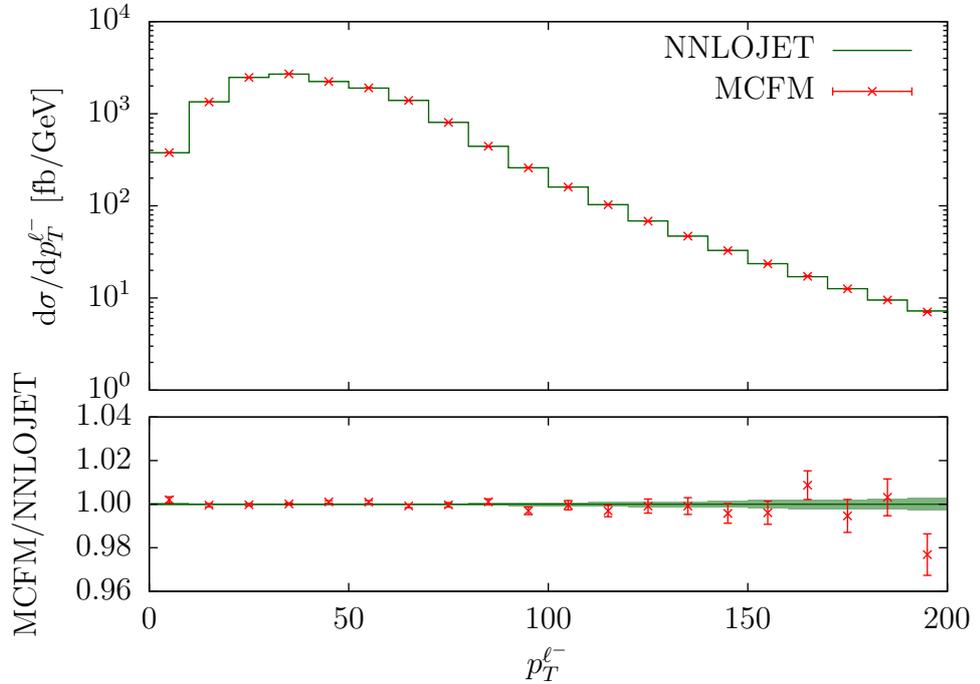


Figure B.5: The transverse momentum distribution of the negatively charged lepton in inclusive Z +jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical.

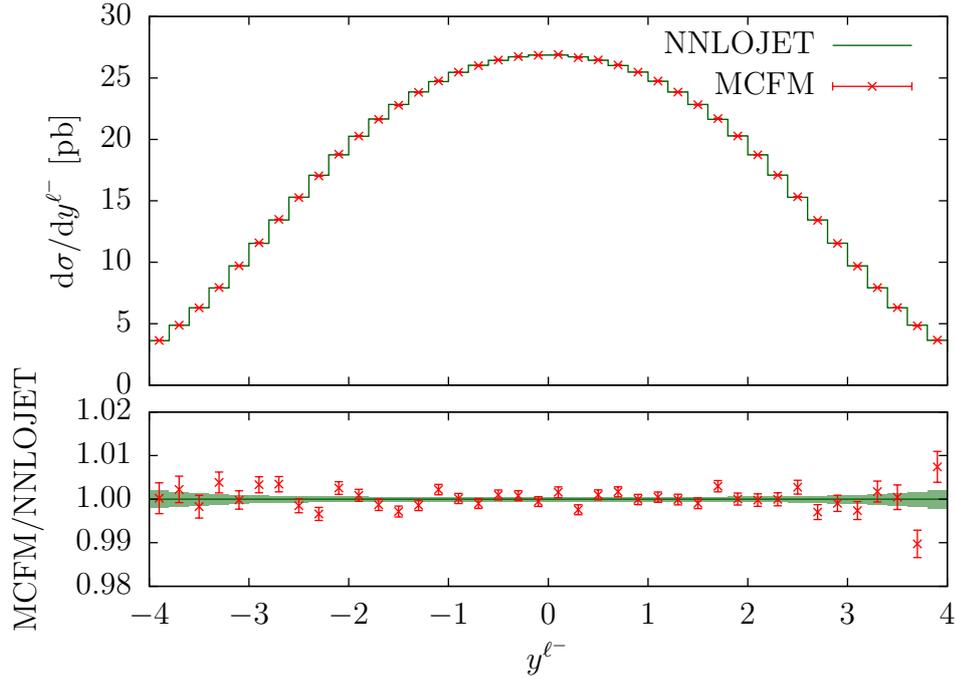


Figure B.6: The rapidity distribution of the negatively charged lepton in inclusive Z +jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical.

phase spaces are also validated using this calculation.

This calculation also partially tests the double real and real virtual subtraction terms for Z +jet at NNLO. The real subtraction term for Z +2 jet at NLO is the first block of terms in the double real Z +jet NNLO subtraction term, $d\hat{\sigma}_{ij}^{S,a}$, described in section 2.4.1. Similarly the virtual subtraction term for Z +2 at NLO is the first block of terms in the real virtual Z +jet NNLO subtraction term, $d\hat{\sigma}_{ij}^{T,a}$, described in section 2.4.2.

The NLO calculation for Z +2 jet production using the NNLOJET code gives the total cross section,

$$\sigma_{\text{NLO}}^{ZJJ} = (33.903 \pm 0.042) \text{ pb},$$

where the uncertainty is the statistical uncertainty from the numerical integration. Performing the same calculation using MCFM gives,

$$\sigma_{\text{NLO}}^{ZJJ} = (33.910 \pm 0.034) \text{ pb}.$$

We can conclude that there is agreement between the two calculations up to per-mille

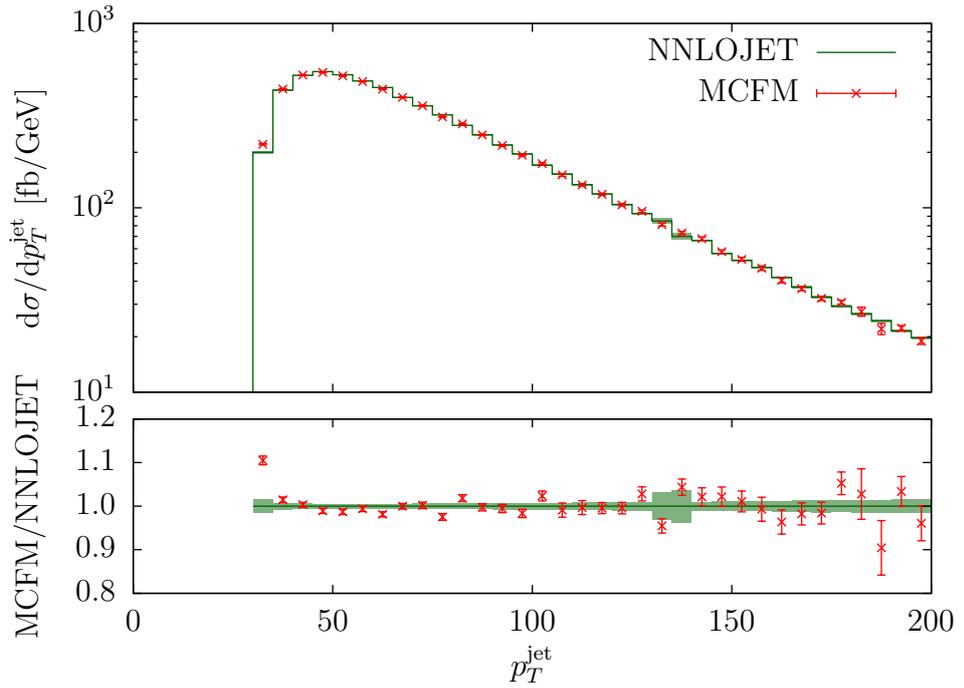


Figure B.7: The transverse momentum distribution of the leading jet in inclusive $Z+2$ jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical.

accuracy for the total cross section. Fig. B.7 shows the leading jet transverse momentum distribution for $Z+2$ jet production at NLO calculated using both NNLOJET (green) and MCFM (red). Here we observe a discrepancy between the two calculations in the 30 GeV bin. This can be understood because in this bin the leading and sub-leading jet are both close to the jet cut threshold. This leads to a numerically unstable region of phase space where it is possible that either or both calculations are underestimating the statistical uncertainties associated with this bin. This problematic region of phase space is not an issue for the full Z +jet at NNLO calculation.

Other than the bin close to the jet cut, Fig. B.7 shows per-cent agreement between the NNLOJET and MCFM implementations for the $Z+2$ jet production at NLO.

Fig. B.8 shows the rapidity distribution of the leading jet in inclusive $Z+2$ jet production at NLO. Per-cent level agreement is observed between the two implementations of the calculation.

Fig. B.9 shows the transverse momentum distribution of the Z boson in inclusive $Z+2$ jet production at NLO. Note that there is no Sudakov shoulder phenomenon because the leading order contribution has two hard jets for the Z to recoil against at

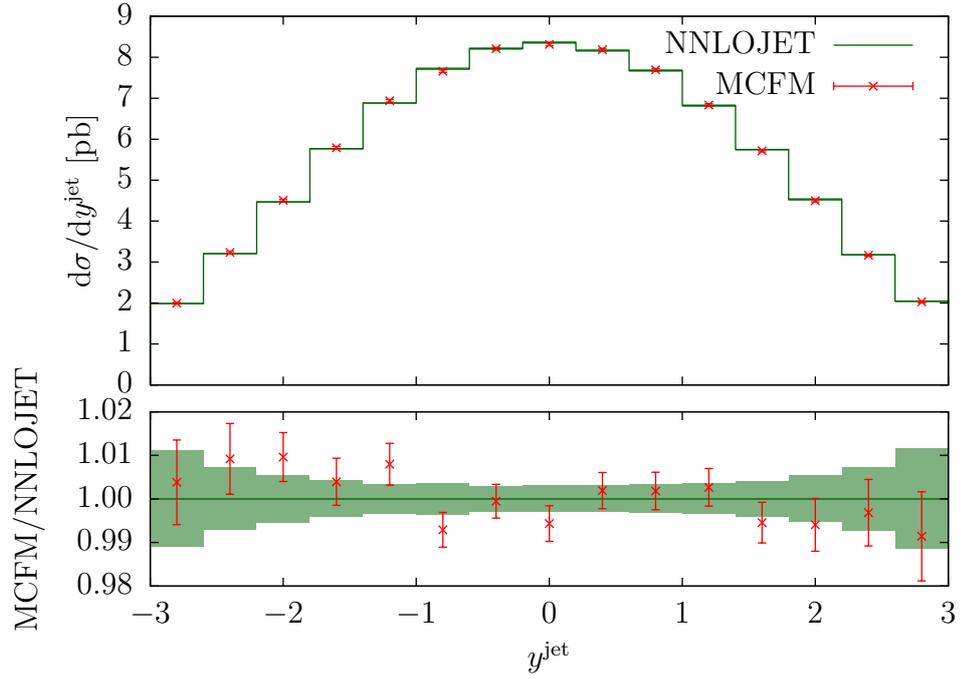


Figure B.8: The rapidity distribution of the leading jet in inclusive $Z+2$ jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical.

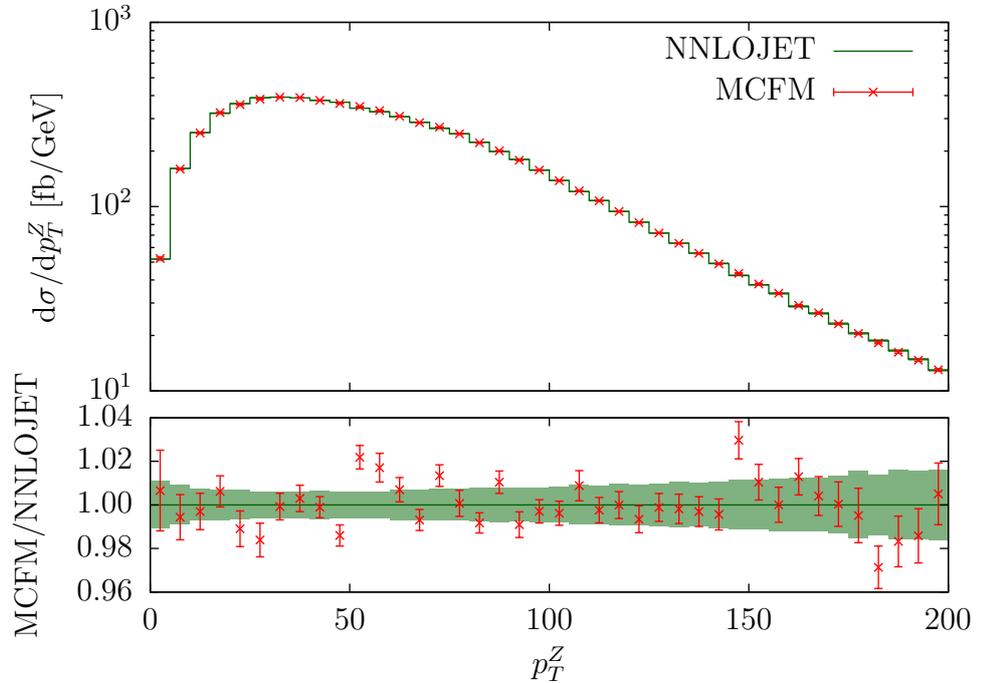


Figure B.9: The transverse momentum distribution of the Z boson in inclusive $Z+2$ jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical.

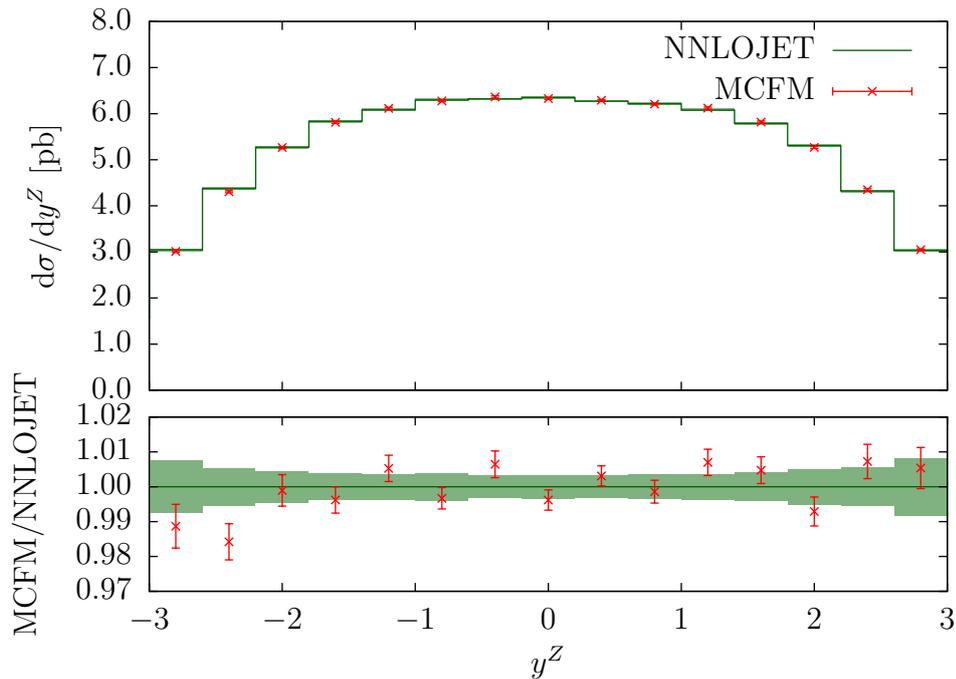


Figure B.10: The rapidity distribution of the Z boson in inclusive $Z+2$ jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical.

low transverse momentum. Agreement is observed between the NNLOJET and MCFM implementations of the calculation within 4% on each bin.

Fig. B.10 shows the rapidity distribution of the Z boson in inclusive $Z+2$ jet production at NLO. Again excellent agreement is observed between the two implementations, where each bin is within agreement within 2%.

Fig. B.11 shows the transverse momentum distribution of the negatively charged lepton in inclusive $Z+2$ jet production at NLO. Here agreement is observed between the two calculations up to 4% on each bin.

Finally, Fig. B.12 shows the rapidity distribution of the negatively charged lepton in inclusive $Z+2$ jet production at NLO. Here per-cent level agreement is observed between the two calculations.

B.2 Technical cut plots

As discussed in section 4.5, an important validation of any subtraction term involving unresolved real emissions is the dependence of the integral on the technical cut pa-

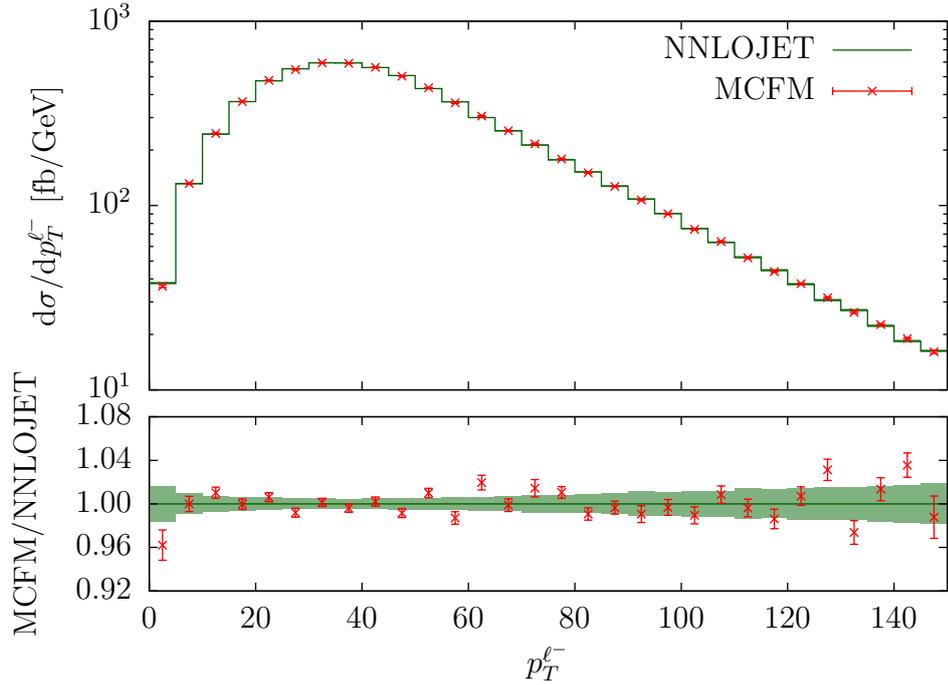


Figure B.11: The transverse momentum distribution of the negatively charged lepton in inclusive $Z+2$ jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical.

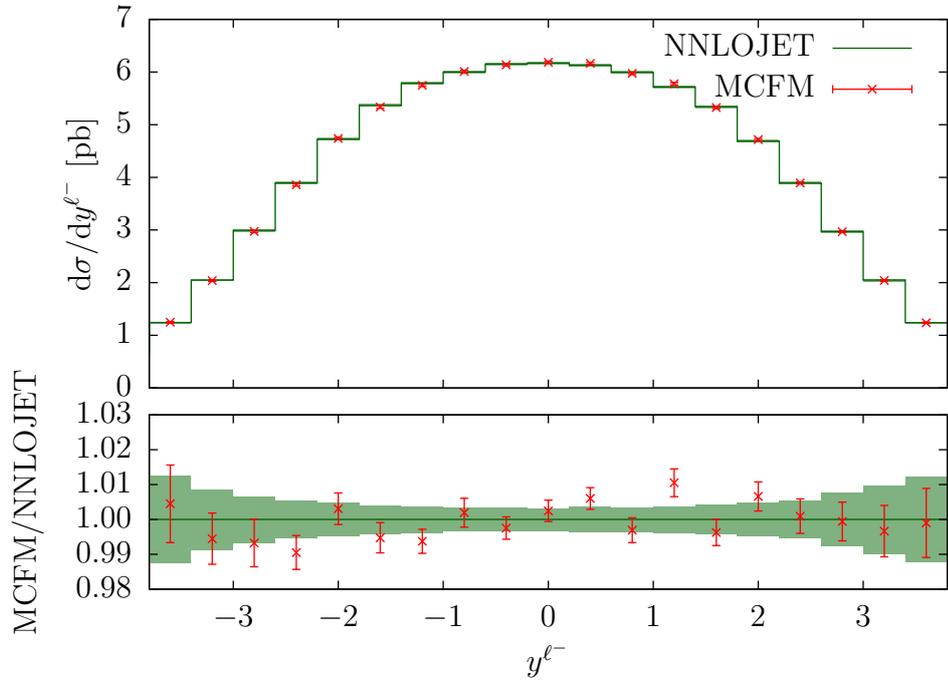


Figure B.12: The rapidity distribution of the negatively charged lepton in inclusive $Z+2$ jet production at NLO. The green band denotes the calculation using the NNLOJET code and the red data points are an equivalent calculation using MCFM with the same kinematical cuts, scale and PDF choices. The uncertainties shown are statistical.

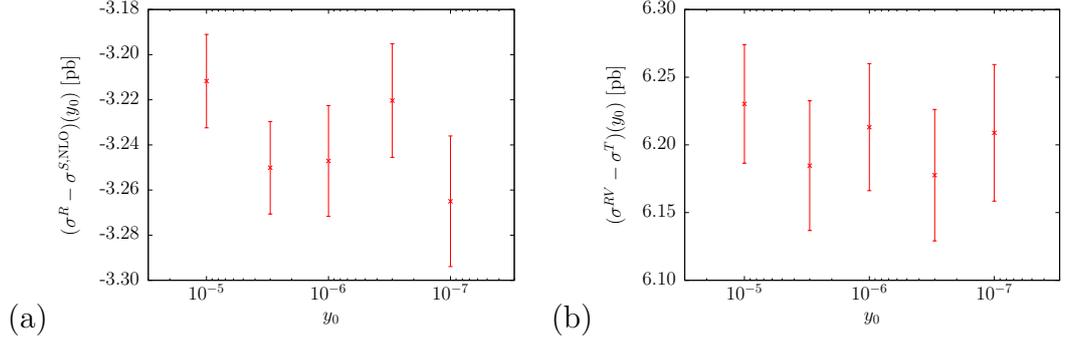


Figure B.13: Contributions to the cross section for Z +jet production for different values of the technical cut, y_0 , for (a) real corrections at NLO and (b) real virtual corrections at NNLO. The error bars on the datapoints are statistical.

parameter. These distributions were shown for the double real phase space integrals in Fig. 4.5.

Fig. B.13 demonstrates similar plots for the single unresolved emissions coming from the both the real corrections at NLO (Fig. B.13(a)) and the real virtual corrections at NNLO (Fig. B.13(b)). It is clear that there is no residual technical cut dependence and the subtraction terms in both cases are performing well. For all of the other distributions presented in this thesis, y_0 was chosen to be $y_0 = 10^{-6}$ for both the real and real virtual corrections to the cross section. The comparison of Fig. 4.5 to Fig. B.13 demonstrates that code is significantly more stable when considering single unresolved emissions compared to double unresolved emissions.

Whilst it is informative to check the technical cut dependence of the total cross section, it is not a complete validation of the real phase space integration procedure. For observables where the weight is a sharply falling function, for example in transverse momentum distributions, the technical cut dependence on the total cross section is only representative of the first few bins for the observable. Indeed, it is important that we can demonstrate that there is no residual technical cut dependence in bins which form a negligible contribution to the total cross section.

Fig. B.14 shows the transverse momentum distribution of the leading jet, double real region a contribution normalised to the high statistics double real region a result used in section 6.1 for various choices of y_0 . Whilst the statistical uncertainties are large, it is clear that there is largely good agreement between the high statistics result and the runs using a range of technical cuts. There is no net trend in the results and the

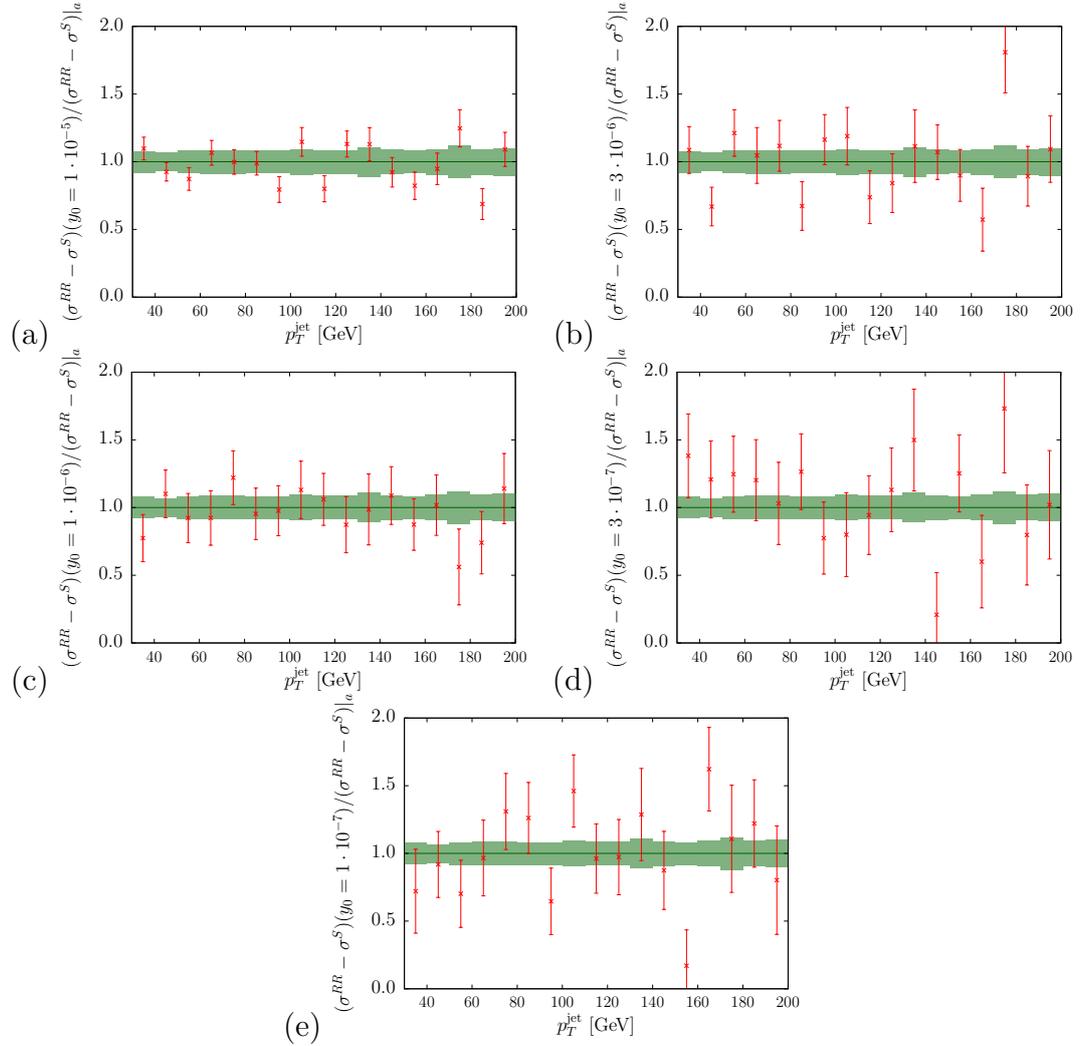


Figure B.14: The transverse momentum distribution for the leading jet, double real region a contribution normalised to the high statistics double real region a result used in section 6.1. The errors on each datapoint are statistical. The technical cut values are (a) $y_0 = 1 \cdot 10^{-5}$, (b) $y_0 = 3 \cdot 10^{-6}$, (c) $y_0 = 1 \cdot 10^{-6}$, (d), $y_0 = 3 \cdot 10^{-7}$, (e) $y_0 = 1 \cdot 10^{-7}$. The green band denotes the statistical uncertainty on the high statistics calculation.

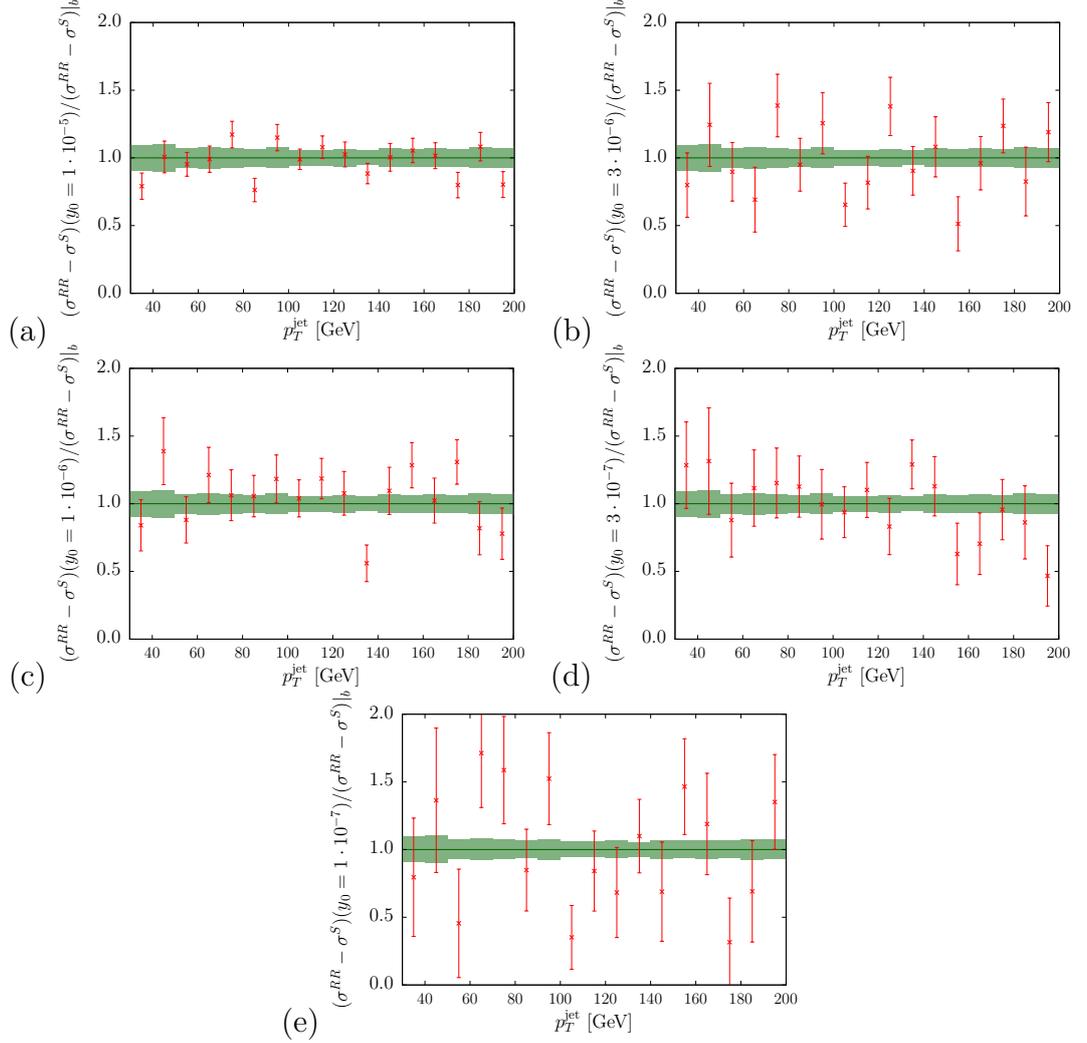


Figure B.15: The transverse momentum distribution for the leading jet, double real region b contribution normalised to the high statistics double real region b result used in section 6.1. The errors on each datapoint are statistical. The technical cut values are (a) $y_0 = 1 \cdot 10^{-5}$, (b) $y_0 = 3 \cdot 10^{-6}$, (c) $y_0 = 1 \cdot 10^{-6}$, (d), $y_0 = 3 \cdot 10^{-7}$, (e) $y_0 = 1 \cdot 10^{-7}$. The green band denotes the statistical uncertainty from the high statistics calculation.

fluctuations are entirely statistical. Similar conclusions can be drawn from Fig. B.15, which shows equivalent distributions for the double real region b calculation.

A similar exercise can be performed for checking the dependence of the calculation on the dynamic reduced technical cut, controlled by the parameter δ_0 . Fig. B.16 shows the transverse momentum distribution double real region a contribution normalised to the high statistics calculation for various choices of δ_0 . It is clear that not only is there excellent agreement for all the values of δ_0 considered, but also the statistical errors in each bin have been reduced significantly in comparison to there being no dynamic cut applied. Indeed, the size of the error bars is comparable to the high statistics

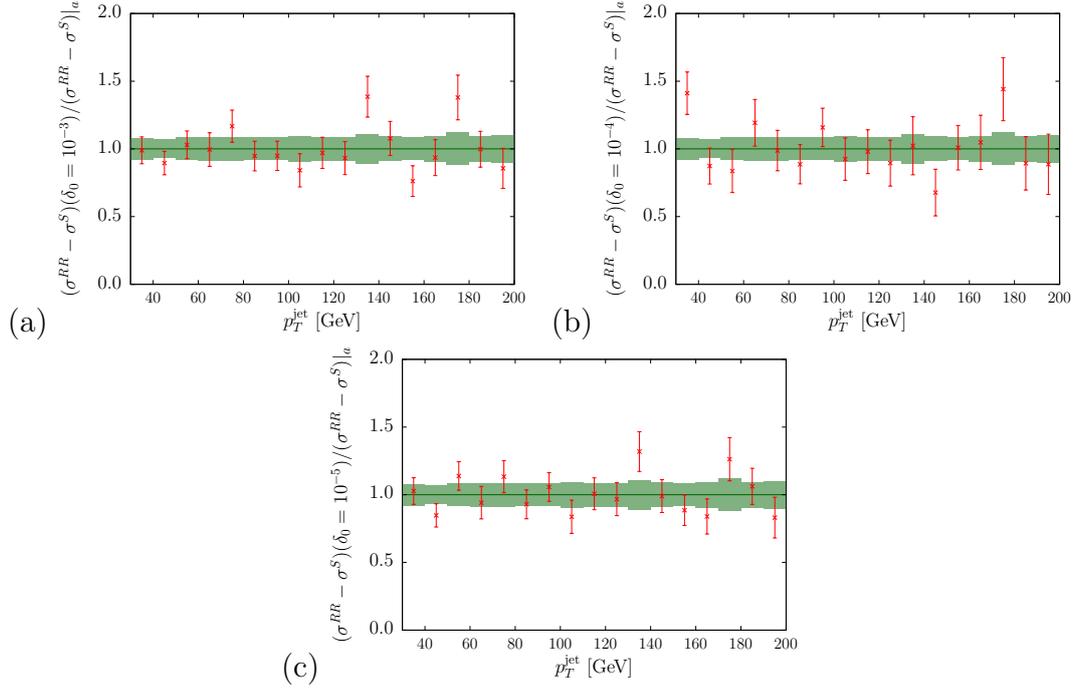


Figure B.16: The transverse momentum distribution for the leading jet, double real region a contribution normalised to the high statistics double real region a result used in section 6.1. The errors on each datapoint are statistical. The reduced technical cut values are (a) $\delta_0 = 10^{-3}$, (b) $\delta_0 = 10^{-4}$, (c) $\delta_0 = 10^{-5}$. The green band denotes the statistical uncertainty from the high statistics calculation.

calculation for all of the choices of δ_0 . This shows that the dynamic cut offers a huge performance upgrade in terms of the speed of the convergence of the calculation for differential distributions. Again, similar conclusions can be drawn from Fig. B.17, which shows the same distributions for the double real region b integral.

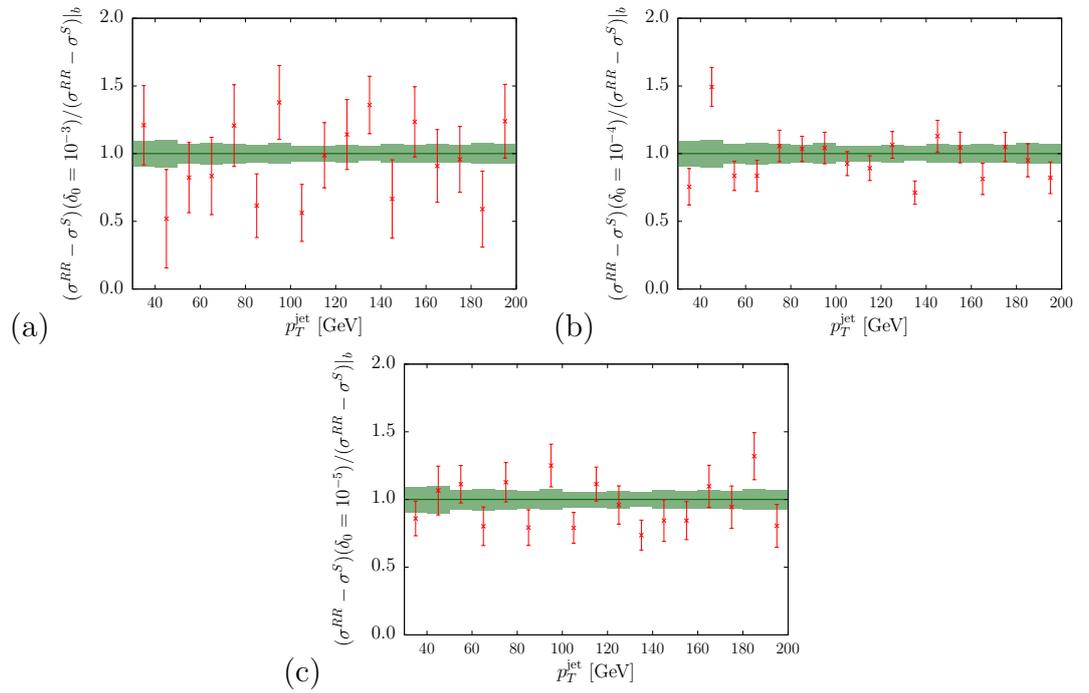


Figure B.17: The transverse momentum distribution for the leading jet, double real region b contribution normalised to the high statistics double real region b result used in section 6.1. The errors on each datapoint are statistical. The reduced technical cut values are (a) $\delta_0 = 10^{-3}$, (b) $\delta_0 = 10^{-4}$, (c) $\delta_0 = 10^{-5}$. The green band denotes the statistical uncertainty from the high statistics calculation.