# Generators and the (Accelerated) Future

#### J Isaacson

Particle Theory Department, Fermi National Accelerator Laboratory, Batavia, IL 60510, USA E-mail: isaacson@fnal.gov

**Abstract.** With the High Luminosity LHC coming online in the near future, event generators will need to provide very large event samples to match the experimental precision. Currently, the estimated cost to generate these events exceeds the computing budget of the LHC experiments. To address these issues, the computing efficiency of event generators need to be improved. Many different approaches are being taken to achieve this goal. I will cover the ongoing work on implementing event generators on the GPUs, machine learning the matrix element, machine learning the phase space, and minimizing the number of negative weight events.

#### 1. Introduction

Collider experiments are developed by the high energy physics community to understand the universe at the smallest distances. Event generators are tools that have been developed to connect the theory calculations to the experimental measurements. This connection is made through simulating the collisions event-by-event through the use of Monte-Carlo methods. The event generators can be broken down into separate components describing the physics at different scales involved in the collision [1, 2]. Starting from the highest scale and working to the lowest scale, these consist of the hard matrix element, parton showers, and hadronization. Our fundamental theoretical understanding is encoded in the hard matrix element through matrix element generators. Matrix element generators are able to compute processes at tree-level [3–10], one-loop level [11–18], and can handle a broad range of Beyond the Standard Model scenarios in an automated fashion [19–21].

As typically is the case, increasing the flexibility of these programs comes at a cost of computational efficiency. The alternative to these automated tools is to obtain analytic calculations for each process. While this is feasible for low multiplicity, there are no analytic results for high multiplicity processes. Therefore, the trade-off in efficiency is acceptable in order to simulate the needed higher multiplicity events. High multiplicity events will become even more important with the high-luminosity phase of the Large Hadron Collider (LHC) starting up in the near future. The datasets collected in this phase of the LHC will have an unprecedented number of events. Ensuring that the event generator statistics are not the dominant uncertainty in these analyses requires a similar sized dataset. However with current techniques, this would far exceed the computational budget for the experiments (see the left panel of Fig. 1). The cost of these events are dominated by the hard matrix element calculation, which scales exponentially as a function of the multiplicity (see the right panel of Fig. 1).

To address these computing limitations, many different solutions are being investigated. Here we will focus on four major categories of advancements being studied. The first is developing



Figure 1. Left: Estimated computing requirements for current (open circles), conservative speedups (closed vertical triangles), and aggressive speedups (closed inverted triangles) as a function of time. Additionally, estimates of the expected computing budget are given as solid black lines. The figure is reproduced from Ref. [22]. Right: CPU hours required to generate one million events as a function of number of final state jets. The calculation of the hard matrix element is shown in ref, and the calculation of the parton shower is shown in green and blue. This figure is reproduced from Ref. [23].

matrix element generators that can take advantage of hardware accelerators such as graphics processing units (GPUs), and will be discussed in Sec. 2. The second is using machine learning to develop a quick but accurate estimate of the matrix element to reduce the need to perform the full calculation. Details of this approach can be found in Sec. 3. Thirdly, machine learning can also be used to more efficiently sample the phase space integrals involved in simulating the events for the colliders, which is explained further in Sec. 4. Finally, when going beyond tree-level descriptions, negative weights are unavoidable. Negative weights drastically decrease the effective sample size. Methods to reduce the negative weight fraction are reviewed in Sec. 5.

#### 2. GPU Matrix Element Generators

There are three major efforts investigating the use of GPUs for matrix element calculations. Each group takes a different approach in handling the conversion from CPU codes to GPU codes. These approaches are the ones developed in BlockGen [24], MadFlow [25], and MadGraph-GPU [26].

#### 2.1. BlockGen

The BlockGen framework [24] implements the Berends-Giele recursion relation [27, 28] directly onto GPUs. They performed a comprehensive study comparing color-ordered (BlockGen-CO<sub> $\Sigma$ </sub>) and color-dressed (BlockGen-CD<sub>MC</sub>) implementations. The authors found that both algorithms are memory bound and not compute bound on the GPUs. The memory required by each algorithm can be found in Fig. 2. The large memory footprint prevents all the needed memory fitting in the local memory, resulting in memory fetching dominating the time per event. However, significant improvement over traditional CPU implementations was found (see Fig 3).

#### 2.2. MadFlow

The MadFlow framework [25] develops an interface with the MadGraph5\_aMC@NLO program [17] to generate Tensorflow [29] code. Since Tensorflow implements GPU calculations,



**Figure 2.** <u>Memory scaling for typical concurrent partonic calculations as a function of number</u> of particles.<sup>3</sup> The top right panel shows memory independent of the number of threads, while the left and bottom right panels show the additional memory per thread. Figure is reproduced from Ref. [24].





Figure 3. The comparison of different matrix element generator implementations. The crosses denoted GPU codes, while the circles denote CPU codes. The solid black (blue) line denotes the best performance amongst all implementations on GPUs (CPUs). Figure is reproduced from Ref. [24].

the generated code from MadFlow can be directly run on either CPUs or GPUs. The MadFlow authors compared the runtime of their implementation on a variety of GPUs and CPUs to determine the overall performance improvement. While there is drastic improvements in computation speed at low multiplicities, the gains at high multiplicity are less significant (see Fig. 4).

#### 2.3. MadGraph-GPU

The MadGraph-GPU framework builds on the MadGraph5\_aMC@NLO program to autogenerate CUDA code [26]. This was handled through a recursive process of optimizing the output of the C++ generator line by line, instead of redesigning the code from the ground up for GPUs. While only preliminary results are currently available for the process  $e^+e^- \rightarrow \mu^+\mu^-$ , the performance is very promising (see Tab. 1).

#### 3. Machine Learning: Matrix Elements

In addition to trying to more efficiently use the available resources on high-performance computers, other groups are investigating methods of estimating the matrix element efficiently. These estimates can then be used to reduce the number of exact matrix evaluations required through a multiple step unweighting procedure. This new unweighting technique is referred to as Neural Rejection Sampling [30]. The proposed algorithm first uses an approximation for the matrix element, and then does a first accept-reject step. If the event fails, then there is no need



**Figure 4.** Performance of the MadFlow algorithm for low multiplicity (left) and high multiplicity (right). The performance of various GPUs are shown in blue, and the performance of various CPUs is shown in red. Figure is reproduced from Ref. [25].

**Table 1.** Preliminary performance results for the MadGraph-GPU code for the process  $e^+e^- \rightarrow \mu^+\mu^-$ . Values in the table are reproduced from Ref. [26].

Implementation	MEs/sec (double precision)	Ratio to C++ code
MadEvent Fortran	$1.50 \times 10^6$	1.15
Standalone C++	$1.31 \times 10^6$	1.00
Standalone CUDA: NVidia V100	$1.37 \times 10^9$	1050
Standalone CUDA: NVidia T4	$4.01 \times 10^{7}$	31

to evaluate the expensive exact matrix element. However, if the event is accepted, the exact matrix element is evaluated and another accept-reject step is performed to obtain an unweighted event. The efficiency of this method can be found in Tab. 2.

**Table 2.** Performance measures for unweighting W + 4j events at the LHC.  $\epsilon_{\text{full}}$  is the overall efficiency,  $\epsilon_{(1\text{st}/2\text{nd}),\text{surr}}$  is the efficiency for the surrogate and full matrix element respectively, and  $\langle t_{\text{full}} \rangle / \langle t_{\text{surr}} \rangle$  is the ratio of time for the full matrix element to the surrogate. Further details of each parameter and data can be found in Ref. [30].

Process	$\epsilon_{\mathrm{full}}$	$\langle t_{\rm full} \rangle / \langle t_{\rm surr} \rangle$	$\epsilon_{\rm 1st,surr}$	$\epsilon^{med}_{\rm 2nd, surr}$
$\begin{array}{c} dg \rightarrow e^- \bar{\nu}_e gggu \\ dd \rightarrow e^- \bar{\nu}_e ggdu \\ ud \rightarrow e^- \bar{\nu}_e duu\bar{d} \end{array}$	$\begin{array}{c} 1.4 \times 10^{-3} \\ 3.1 \times 10^{-4} \\ 3.6 \times 10^{-4} \end{array}$	$667 \\ 162 \\ 25$	$\begin{array}{c} 7.1 \times 10^{-4} \\ 1.1 \times 10^{-4} \\ 1.3 \times 10^{-4} \end{array}$	$\begin{array}{c} 5.3\times 10^{-2} \\ 8.5\times 10^{-2} \\ 7.3\times 10^{-2} \end{array}$

To obtain optimal performance of the neural rejection algorithm requires an extremely accurate matrix element estimate. This can be obtained with the help of a factorization aware Matrix element emulator. Such an emulator has been proposed in Ref. [31]. It uses information



#### events [33–42], and generating the phase space efficiently [33, 43–48]. Dipole NN'single Dipole NN 'ensemble' [2002.07516] single [2002.07516] ensemble

## 4.1. From Events

The generation of events from a set of pre-generated events are able to quickly generate additional events that are similar to the existing ones. However, these approaches have no guarantees on maintaining the total cross section and struggle with peaked structures in the integrand. The work in this direction mainly focuses on the use of GANs to generate the events, and is very successful in reproducing smooth distributions. The current state of the art for these methods can be seen in Fig. 6.

# 4.2. Phase Space

The other approach attempts to use machine learning to greatly increase the unweighting efficiency. The advantage of this approach is that it does not require a pre-generated sample and through the unweighting procedure guarantees to reproduce the correct cross section. The downside of this approach is that the networks are more computationally intensive, and still requires an unweighting step. These two issues lead to a drastically slower method of generating events than the above approach. The main architecture investigated are invertible networks, and some promising results have been found in Refs. [44–46] (see the right panel of Fig. 6 and Tab. 3).



**Figure 6.** Left and Center: Event generation through the use of a trained GAN network. The left plot shows the accuracy of GANs for smooth distributions, while the center plot shows the accuracy of these techniques for peaked distributions. The different curves in the center plot are modifications added to the loss to improve the accuracy. Right: The weight distribution for even generation. The unweighting efficiency is directly related to the narrowness of the distribution. The neural network (green) shows a drastic improvement over the previous optimal method VEGAS (orange). Left and center figures are reproduced from Ref. [37], and the right figure is reproduced from Ref. [44].

unweighting $\epsilon$	efficiency	$\begin{array}{l} \text{LO QCD} \\ n = 0 \end{array}$	n = 1	n = 2	n = 3	n = 4
$W^+ + n$ jets	Sherpa NN+NF Gain	$2.8 \times 10^{-1}$ $6.1 \times 10^{-1}$ 2.2	$3.8 \times 10^{-2}$ $1.2 \times 10^{-1}$ 3.3	$7.5 \times 10^{-3}$ $1.0 \times 10^{-2}$ 1.4	$1.5 \times 10^{-3}$ $1.8 \times 10^{-3}$ 1.2	$8.3 \times 10^{-4}$ $8.9 \times 10^{-4}$ 1.1
$W^- + n$ jets	Sherpa NN+NF Gain	$2.9 \times 10^{-1}$ $7.0 \times 10^{-1}$ 2.4	$\begin{array}{c} 4.0\times 10^{-2} \\ 1.5\times 10^{-1} \\ 3.3 \end{array}$	$7.7 \times 10^{-3} \\ 1.1 \times 10^{-2} \\ 1.4$	$2.0 \times 10^{-3}$ $2.2 \times 10^{-3}$ 1.1	$9.7 \times 10^{-4}$ $7.9 \times 10^{-4}$ 0.82
$Z^- + n$ jets	Sherpa NN+NF Gain	$3.1 \times 10^{-1}$ $3.8 \times 10^{-1}$ 1.2	$3.6 \times 10^{-2}$ $1.0 \times 10^{-1}$ 2.9	$1.5 \times 10^{-2}$ $1.4 \times 10^{-2}$ 0.91	$4.7 \times 10^{-3}$ $2.4 \times 10^{-3}$ 0.51	

**Table 3.** Unweighting efficiencies for the default Sherpa and the improvements from neural network phase space generation for a variety of processes of interest. Data for the table is taken from Ref. [46].

#### 5. Reducing Negative Weights

In higher order calculations, negative event weights are unavoidable. Negative weights arise in the color dipole terms, overestimation of real-emission matrix elements, and in matching to parton showers. The impact of negative weight events is a reduction in effective sample size, and thus statistical accuracy. Given a set of events with weights either positive one or negative one, with a fraction of  $\epsilon$  events with negative weight, the additional events required for the same accuracy is given as:

$$f(\epsilon) = \frac{1}{(1 - 2\epsilon)^2}.$$
(1)

This quickly increases as  $\epsilon$  approaches 50%. Therefore, to improve the efficiency of event generation, the fraction of negative weight events needs to be as small as possible. There are two methods of reducing the negative weight fraction.

The first is through the use of resampling the events to produce only positive weight events. This can be accomplished through the use of a positiver resampler [49], a neural network resampler [50], or cell resampling [51].

In the positiver resampler [49], events are rescaled by  $\frac{\sum_i w_i}{\sum_i |w_i|}$ , partially unweighted, and then rescaled to ensure that the original histogram bin height is restored. The reduction in events required to reach a given precision can be seen in Tab. 4.

**Table 4.** Number of weighted events generated and number of events that pass analysis cuts from merged W + 0j, 1j, 2j at next-to-leading order calculation. Data is taken from Ref. [49].

Sample	Total number of events	Events included in analysis
weighted	5.3M	195k
positive only	$3.2\mathrm{M}$	121k
unweighted	$1.5\mathrm{M}$	52k
Positive $\operatorname{Resampler}(t)$	659k	25k
Positive Resampler $(p_T^W)$	33k	33k



Figure 7. Reduction in required number of events generated through a neural resampling procedure. Figure is reproduced from Ref. [50].

In the neural resampler [50], a weight and variance are obtained through a neural network to estimate how close the result is to uniformly weighted sample. The weight and variance are then used to calculate a scaling factor. This scaling factor is used to perform an accept-reject step and then the event weight is increased by this scaling factor. One advantage to this technique is that the variance of the original sample is maintained. The validation and reduction in events generated for  $pp \rightarrow t\bar{t}$  at NLO is shown in Fig. 7.

In the cell resampler [51], a phase space measure is defined to calculate the distance between two points. The negative weight events are then chosen as a seed of a cell. The cell size is increased until the sum of the weights in the cell is at least twice the negative weight seed. Afterwards, all the events in the cell are scaled  $\frac{\sum_{i \in \mathcal{C}} w_i}{\sum_{i \in \mathcal{C}} |w_i|}$ . If there is an insufficient number of



**Figure 8.** Left: Validation that the cell resampler reproduces the expected result. Right: The fraction of positive weight and negative weight events before and after the cell resampler. The figure is reproduced from Ref. [51].

events in a local region to make the weights positive, the negative weight is kept. This drastically decreases the number of events required and the fraction of negative weights (see Fig. 8).

Finally, in addition to resampling techniques, the algorithms can be improved to reduce the negative weight fraction. In Ref. [52], three improvements were proposed. These improvements were matching to leading color showers, shower vetos on hard emission events, and local K-factors from the core process. The combination of these three improvements halved the negative weight fraction within the study. The breakdown of the effect of each improvement can be found in Tab. 5.

Table 5	i. Impi	rovements	made in t	the Sherpa	event	generator	to reduce	e the	fraction	of nega	tive
weight e	events.	Data is ta	ken from	Ref. [52].					_		

	Negative Weight Fraction
Default	18.1%
Leading Color Mode + shower veto on $\mathbb{H}$ -events + local K-factor from core	14.0% 9.6% 9.1%

### 6. Conclusions

In order to match the statistical precision of the high-luminosity LHC, a large number of events need to be generated with the help of MC programs. The currently available event generators require a computing budget that exceeds the provided budget to reach the desired goal. Therefore, there have been many steps taken to improve event generators to reduce the cost per event.

Speed improvements through the use of dedicated GPU implementations for the hard matrix element show promise in reducing the cost, but may not be sufficient on their own. Using machine learning techniques to provide faster unweighting methods and improved phase space generation also help to increase the overall speed of event generation. Finally, studying negative weight reduction techniques works on minimizing number of simulated events required. While each of these steps help to work towards reducing the computing costs to the desired level, there is still more that needs to be done in order to meet the experimental requirements.

#### Acknowledgments

I would like to thank Stefan Höche for useful discussions and advice on the preparation of these proceedings. This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

#### References

- [1] Webber B 1986 Ann. Rev. Nucl. Part. Sci. 36 253-286 URL http://inspirebeta.net/ record/18125
- [2] Buckley A et al. 2011 Phys. Rept. 504 145-233 (Preprint 1101.2599) URL http:// inspirebeta.net/record/884202
- Kanaki A and Papadopoulos C G 2000 Comput. Phys. Commun. 132 306-315 (Preprint hep-ph/0002082) URL http://inspirehep.net/search?p=hep-ph/0002082
- [4] Papadopoulos C G 2001 Comput. Phys. Commun. 137 247-254 (Preprint hep-ph/0007335) URL http://inspirehep.net/search?p=hep-ph/0007335
- [5] Krauss F, Kuhn R and Soff G 2002 JHEP 02 044 (Preprint hep-ph/0109036) URL http://inspirehep.net/search?p=hep-ph/0109036
- [6] Gleisberg T and Höche S 2008 JHEP 12 039 (Preprint 0808.3674) URL http:// inspirehep.net/record/793879
- [7] Maltoni F and Stelzer T 2003 JHEP 02 027 (Preprint hep-ph/0208156) URL http: //inspirehep.net/search?p=hep-ph/0208156
- [8] Alwall J et al. 2007 JHEP 09 028 (Preprint 0706.2334) URL http://inspirehep.net/ search?p=arXiv:0706.2334
- [9] Alwall J, Herquet M, Maltoni F, Mattelaer O and Stelzer T 2011 JHEP 06 128 (Preprint 1106.0522) URL http://inspirebeta.net/record/912611
- [10] Mangano M L, Moretti M, Piccinini F, Pittau R and Polosa A D 2003 JHEP 07 001 (Preprint hep-ph/0206293) URL http://inspirehep.net/search?p=hep-ph/0206293
- [11] Ossola G, Papadopoulos C G and Pittau R 2008 JHEP 0803 042 (Preprint 0711.3596)
- [12] Gleisberg T and Krauss F 2008 Eur. Phys. J. C53 501-523 (Preprint 0709.2881) URL http://inspirehep.net/search?p=arXiv:0709.2881
- [13] Berger C F, Bern Z, Dixon L J, Febres-Cordero F, Forde D, Ita H, Kosower D A and Maître D 2008 Phys. Rev. D78 036003 (Preprint 0803.4180) URL http://inspirebeta. net/record/782271
- [14] Bevilacqua G, Czakon M, Garzelli M, van Hameren A, Kardos A et al. 2013 Comput.Phys.Commun. 184 986–997 (Preprint 1110.1499)
- [15] Cascioli F, Maierhöfer P and Pozzorini S 2012 Phys. Rev. Lett. 108 111601 (Preprint 1111.5206) URL http://inspirehep.net/record/946998
- [16] Cullen G, van Deurzen H, Greiner N, Heinrich G, Luisoni G et al. 2014 Eur.Phys.J. C74 3001 (Preprint 1404.7096) URL http://inspirehep.net/record/1292822
- [17] Alwall J, Frederix R, Frixione S, Hirschi V, Maltoni F, Mattelaer O, Shao H S, Stelzer T, Torrielli P and Zaro M 2014 JHEP 07 079 (Preprint 1405.0301)
- [18] Actis S, Denner A, Hofer L, Lang J N, Scharf A and Uccirati S 2017 Comput. Phys. Commun. 214 140–173 (Preprint 1605.01090)

- [19] Christensen N D, de Aquino P, Degrande C, Duhr C, Fuks B, Herquet M, Maltoni F and Schumann S 2011 Eur. Phys. J. C71 1541 (Preprint 0906.2474) URL http: //inspirebeta.net/record/823106
- [20] Degrande C, Duhr C, Fuks B, Grellscheid D, Mattelaer O and Reiter T 2012 Comput. Phys. Commun. 183 1201–1214 (Preprint 1108.2040)
- [21] Staub F 2014 Comput. Phys. Commun. 185 1773–1790 (Preprint 1309.7223)
- [22] Calafiura P, Catmore J, Costanzo D and Di Girolamo A 2020 ATLAS HL-LHC Computing Conceptual Design Report Tech. Rep. CERN-LHCC-2020-015. LHCC-G-178 CERN Geneva URL https://cds.cern.ch/record/2729668
- [23] Höche S, Prestel S and Schulz H 2019 Phys. Rev. D 100 014024 (Preprint 1905.05120)
- [24] Bothmann E, Giele W, Hoeche S, Isaacson J and Knobbe M 2021 (Preprint 2106.06507)
- [25] Carrazza S, Cruz-Martinez J, Rossi M and Zaro M 2021 Eur. Phys. J. C 81 656 (Preprint 2106.10279)
- [26] Valassi A, Roiser S, Mattelaer O and Hageboeck S 2021 EPJ Web Conf. 251 03045 (Preprint 2106.12631)
- [27] Berends F A and Giele W T 1988 Nucl. Phys. B306 759 URL http://inspirehep.net/ search?j=NUPHA,B306,759
- [28] Badger S, Biedermann B, Hackl L, Plefka J, Schuster T and Uwer P 2013 Phys. Rev. D87 034011 (Preprint 1206.2381)
- [29] Abadi M, Agarwal A, Barham P, Brevdo E, Chen Z, Citro C, Corrado G S, Davis A, Dean J, Devin M, Ghemawat S, Goodfellow I, Harp A, Irving G, Isard M, Jia Y, Jozefowicz R, Kaiser L, Kudlur M, Levenberg J, Mané D, Monga R, Moore S, Murray D, Olah C, Schuster M, Shlens J, Steiner B, Sutskever I, Talwar K, Tucker P, Vanhoucke V, Vasudevan V, Viégas F, Vinyals O, Warden P, Wattenberg M, Wicke M, Yu Y and Zheng X 2015 TensorFlow: Large-scale machine learning on heterogeneous systems software available from tensorflow.org URL https://www.tensorflow.org/
- [30] Danziger K, Janßen T, Schumann S and Siegert F 2021 (*Preprint* 2109.11964)
- [31] Maître D and Truong H 2021 JHEP **11** 066 (Preprint **2107.06625**)
- [32] Catani S and Seymour M H 1997 Nucl. Phys. B485 291-419 (Preprint hep-ph/9605323) URL http://inspirehep.net/search?p=hep-ph/9605323
- [33] Bendavid J 2017 (*Preprint* 1707.00028)
- [34] Otten S, Caron S, de Swart W, van Beekveld M, Hendriks L, van Leeuwen C, Podareanu D, Ruiz de Austri R and Verheyen R 2021 Nature Commun. 12 2985 (Preprint 1901.00875)
- [35] Hashemi B, Amin N, Datta K, Olivito D and Pierini M 2019 (Preprint 1901.05282)
- [36] Di Sipio R, Faucci Giannelli M, Ketabchi Haghighat S and Palazzo S 2019 JHEP 08 110 (Preprint 1903.02433)
- [37] Butter A, Plehn T and Winterhalder R 2019 SciPost Phys. 7 075 (Preprint 1907.03764)
- [38] Carrazza S and Dreyer F A 2019 Eur. Phys. J. C 79 979 (Preprint 1909.01359)
- [39] Ahdida C et al. (SHiP) 2019 JINST 14 P11028 (Preprint 1909.04451)
- [40] Butter A, Plehn T and Winterhalder R 2019 (Preprint 1912.08824)
- [41] Butter A, Diefenbacher S, Kasieczka G, Nachman B and Plehn T 2021 SciPost Phys. 10 139 (Preprint 2008.06545)
- [42] Butter A and Plehn T 2020 (*Preprint* 2008.08558)
- [43] Klimek M D and Perelstein M 2020 SciPost Phys. 9 053 (Preprint 1810.11509)

- [44] Bothmann E, Janßen T, Knobbe M, Schmale T and Schumann S 2020 SciPost Phys. 8 069 (Preprint 2001.05478)
- [45] Gao C, Isaacson J and Krause C 2020 Mach. Learn. Sci. Tech. 1 045023 (Preprint 2001.05486)
- [46] Gao C, Höche S, Isaacson J, Krause C and Schulz H 2020 Phys. Rev. D 101 076002 (Preprint 2001.10028)
- [47] Chen I K, Klimek M D and Perelstein M 2021 SciPost Phys. 10 023 (Preprint 2009.07819)
- [48] Stienen B and Verheyen R 2021 SciPost Phys. 10 038 (Preprint 2011.13445)
- [49] Andersen J R, Gütschow C, Maier A and Prestel S 2020 Eur. Phys. J. C 80 1007 (Preprint 2005.09375)
- [50] Nachman B and Thaler J 2020 Phys. Rev. D 102 076004 (Preprint 2007.11586)
- [51] Andersen J R and Maier A 2021 (*Preprint* 2109.07851)
- [52] Danziger K, Höche S and Siegert F 2021 (Preprint 2110.15211)