

Summing bulk quantum numbers with Monte Carlo in spin foam theories

Pietro Donà^{1,2,*} and Pietropaolo Frisoni^{2,†}

¹*Center for Space, Time and the Quantum, 13288 Marseille, France*

²*Department of Physics and Astronomy, University of Western Ontario,
London, Ontario N6A 5B7, Canada*



(Received 7 February 2023; accepted 12 April 2023; published 9 May 2023)

We introduce a strategy to compute EPRL spin foam amplitudes with many internal faces numerically. We work with `sl2cfoam-next`, the state-of-the-art framework to numerically evaluate spin foam transition amplitudes. We find that uniform sampling Monte Carlo is exceptionally effective in approximating the sum over internal quantum numbers of a spin foam amplitude, considerably reducing the computational resources necessary. We apply it to compute large volume divergences of the theory and find surprising numerical evidence that the EPRL vertex renormalization amplitude is instead finite.

DOI: [10.1103/PhysRevD.107.106008](https://doi.org/10.1103/PhysRevD.107.106008)

I. INTRODUCTION

Spin foam theory is the Lorentz covariant version of loop quantum gravity (LQG) and provides a tentative background independent path integral for gravity. It gives dynamics to LQG kinematical states defining transition amplitudes between spin network states [1,2]. The most promising spin foam theory is the EPRL-FK model [3,4]. Various generalizations include the extension to general triangulations [5], the inclusion of a cosmological constant [6,7], and boundary with different signature [8]. These theories promisingly connect with discrete general relativity in the double limit of finer discretization and large areas [9–14].

The field has recently undergone an explosion of numerical methods, providing new tools to address many open questions of the theory. We can compute expectation values and fluctuations of operators in the large spins regime using the complex saddle point analysis and the integration on a Lefschetz thimble using Markov chain Monte Carlo [12,15,16]. It is possible to verify that Regge geometries emerge in the large-scale and small Immirzi parameter regime. A similar result can also be obtained from effective spin foam models [17,18]. Finally, `sl2cfoam` (and its latest iteration `sl2cfoam-next`) is an open source framework to compute EPRL spin foam amplitudes [19,20]. It is based on a divide-and-conquer strategy, and a booster decomposition of the vertex amplitude [21]. The library was already employed to explore the large quantum numbers regime [20,22,23], the infrared divergences of the theory [24,25], the black-to-white hole transition [26], and correlations in the early universe [27]. Very recently, a hybrid approach taking

advantage of all the available techniques was also proposed [28].

This work overcomes one of the principal limitations of `sl2cfoam-next`. The library provides an optimized and efficient framework to compute all the constituents of a spin foam transition amplitude. Nevertheless, numerically computing spin foam amplitudes with many internal faces is prohibitively taxing. There are way too many objects to compute as their number scales exponentially with the number of internal faces. We overcome this problem by evaluating the sums over the internal quantum numbers using statistical frameworks. We explore the possibility of using uniform sampling Monte Carlo and find it surprisingly effective. We can compute amplitudes with slightly better than 1% precision by considering a sampling five orders of magnitude smaller than the total amount of terms of the sum.

We apply this novel technique to compute the melonic self-energy and the vertex renormalization amplitude in the $SU(2)$ BF and EPRL theories. These two amplitudes are the perfect laboratory to test the effectiveness of Monte Carlo as they possess many internal faces. These spin foam amplitudes are believed to be divergent, and their renormalization is crucial to define the continuum limit of the EPRL theory. In the case of the melonic self-energy we find excellent agreement with the numerical results in the literature [24,29], which do not use stochastic methods. This work contains the first computation of the Lorentzian EPRL vertex renormalization amplitude. We surprisingly find numerical evidence for its convergence.

The scripts we use to compute the amplitudes and the notebooks to analyze the data are publicly available at the repository [30]. We perform most of our calculations on the *Narval* cluster of the Digital Research Alliance of Canada.

*pdona@uwo.ca

†pfrisoni@uwo.ca

II. SPIN FOAM TRANSITION AMPLITUDES

We write a spin foam transition amplitude starting from a 2-complex Δ of a simplicial triangulation of the space-time manifold decorated with LQG quantum numbers. Each face is colored with a spin j_f and each edge with an intertwiner i_e .

The spin foam transition amplitude A_Δ is the product of local fundamental amplitudes: a face amplitude $A_f(j_f)$, an edge amplitude $A_e(i_e)$, and a vertex amplitude $A_v(j_f, i_e)$. Finally, we sum over all the possible quantum numbers associated with the bulk of the 2-complex

$$A_\Delta = \sum_{j_f=0}^{\infty} \sum_{i_e} \prod_f A_f(j_f) \prod_e A_e(i_e) \prod_v A_v(j_f, i_e). \quad (1)$$

This work focuses on two spin foam theories: the topological BF $SU(2)$ model and the Lorentzian EPRL model. We introduce them here schematically and report their detailed definition in Appendix A. We refer to reviews [2] or books [1] for a more complete and pedagogical introduction. We use the same notation for the vertex amplitudes in the two models. It is convenient to avoid overburdening the notation and not repeat the same equations twice. We will stress the difference between the two models if necessary. The vertex amplitude for the topological model is defined as

$$A_v(j_f, i_e) = \{15j\}(j_f; i_e), \quad (2)$$

where the $\{15j\}(j_f; i_e)$ is a $SU(2)$ invariant depending on the ten spins and five intertwiners coloring the spin foam vertex. We work with the booster functions decomposition of the Lorentzian EPRL spin foam model introduced in [21]. In this form, the EPRL vertex amplitude is a superposition of $\{15j\}$ symbols weighted by booster functions B_4^γ .

$$A_v(j_f, i_e) = \sum_{l_f=j_f}^{\infty} \sum_{k_e} \{15j\}(j_f, l_f) \times \prod_{e=2}^5 (2k_e + 1) B_4^\gamma(l_f, j_f; i_e, k_e). \quad (3)$$

The presence of the booster functions is the striking difference between the amplitudes (2) and (3). They encode the imposition of the simplicity constraints and the explicit dependence of the theory from the Immirzi parameter γ . They possess a compelling geometrical interpretation of boosted tetrahedra [31]. The edge and face amplitudes are fixed, requiring the correct composition of spin foam amplitudes [32].

$$A_e(i_e) = 2i_e + 1, \quad \text{and} \quad A_f(j_f) = 2j_f + 1. \quad (4)$$

Depending on the details of the 2-complex could be necessary to also multiply by some extra phase and edge-related $SU(2)$ invariants depending on the spin and intertwiners quantum

numbers. They result from our decision to work with a specific recoupling scheme in the vertex amplitudes. We refer to the review [23] for a step-by-step guide on computing them.

We perform all the numerical calculations using `s12cfoam-next`, the state-of-the-art code, to compute spin foam amplitudes with a computer. The library is open source and written in C. It is based on the booster decomposition of the EPRL vertex amplitude, optimizing the available computational resources. We refer to the original paper [20], the review [23], or the book chapter [33] for a detailed description.

One of the main ingredients we mention here is the introduction of a homogeneous truncation parameter Δl to approximate the unbounded convergent sums over the virtual spins l_f in (3).

$$\sum_{l_f=j_f}^{\infty} \rightarrow \sum_{l_f=j_f}^{j_f+\Delta l}. \quad (5)$$

Despite the notation, we emphasize that the truncation parameter Δl is independent of the 2-complex Δ .

III. SUMMING BULK DEGREES OF FREEDOM WITH MONTE CARLO

The library `s12cfoam-next` [20] computes EPRL vertex amplitudes (3) very fast and efficiently.¹ Unfortunately, it is not enough to compute a general spin foam amplitude with many vertices and internal faces. The number of vertex amplitudes we have to calculate, assemble, and sum grows exponentially with the number of bulk faces. We can convince ourselves this is a severe problem with a back-of-the-envelope calculation. Imagine you want to compute an amplitude with F internal faces, and all the spins associated with the inner faces have some characteristic value J . To calculate the amplitude, we must loop through all $(2J+1)^F$ possible values that the internal spins can assume and compute all the vertex amplitudes. Let us assume, optimistically, that we need just 1 μ s of CPU time to obtain them (the actual time is orders of magnitude larger). Suppose we want to calculate an amplitude with 10 internal faces and average spins of order 10. To perform this calculation, we need approximately $21^{10} \mu\text{s} \approx 6$ months of CPU time, which is a lot of time considering our modest requirements and optimistic hypothesis.

We overcome this problem using Monte Carlo to perform the sum over the bulk spins. We rewrite the spin foam amplitude (1) in terms of partial amplitudes

$$A_\Delta = \sum_{j_f=0}^{\infty} a(j_f) \quad \text{with} \quad a(j_f) = \sum_{i_e} \prod_f A_f(j_f) \prod_e A_e(i_e) \prod_v A_v(j_f, i_e). \quad (6)$$

¹It can evaluate the topological BF $SU(2)$ vertex amplitude too.

We include the sum over the bulk intertwiners in the partial amplitudes $a(j_f)$. Those sums are always finite for fixed bulk spins, and we perform them leveraging the tensorial structure of `s12cf Foam-next` without resorting to Monte Carlo methods. We are omitting the evident dependence from the 2-complex Δ of the partial amplitude.

The partial amplitude vanishes if some spins j_f do not satisfy triangular inequalities.² We restrict the sum over the bulk spins in (6) to the set of spins satisfying triangular inequalities that we indicate as \mathcal{I}_Δ . In this way, we eliminate the majority of trivially vanishing partial amplitudes. This step is convenient to obtain an efficient Monte Carlo amplitude estimate.

Generally, \mathcal{I}_Δ is unbounded, therefore is not possible to directly apply Monte Carlo to estimate the amplitude. We circumvent this limitation subdividing \mathcal{I}_Δ into layers \mathcal{J}_k .

$$\mathcal{J}_k = \{j_f \in \mathcal{I}_\Delta \mid \max j_f = k\}. \quad (7)$$

By definition, each layer is finite. Different layers do not overlap $\mathcal{J}_k \cap \mathcal{J}_{k'} = \emptyset$, and the union of all of them forms the original set $\mathcal{I}_\Delta = \cup_{k=0}^{\infty} \mathcal{J}_k$. We reorganize the spin foam amplitude as a sum over layers of layer amplitudes

$$A_\Delta = \sum_{k=0}^{\infty} S_k \quad \text{with} \quad S_k = \sum_{j_f \in \mathcal{J}_k} a(j_f). \quad (8)$$

Each layer amplitude S_k is defined as a sum with a finite number of terms (that, however, grows rapidly with k). We can approximate them with Monte Carlo S_k^{mc} using the procedure described in Appendix B with a fixed amount of samples N_{mc} . In general, the number of layers is infinite.³ We cutoff the sum over the layer to a maximum layer K . This prescription is equivalent to introducing a homogeneous cutoff K to all the bulk spins. The Monte Carlo approximation of a spin foam amplitude with a cutoff K is given by

$$A_\Delta^{mc}(K) = \sum_{k=0}^K S_k^{mc}. \quad (9)$$

The calculation of S_k^{mc} requires a discrete random uniform probability distribution over the layer \mathcal{J}_k . In principle,

²For example, if j_1, j_2, j_3 , and j_4 are the four spins associated with the faces contained in an edge of Δ and

$$\text{Max}(|j_1 - j_2|, |j_3 - j_4|) > \text{Min}(j_1 + j_2, j_3 + j_4) \Rightarrow a_\Delta(j_f) = 0$$

then the set of intertwiners associated with that edge is empty, and the partial amplitude trivially vanishes.

³In some exceptional cases, the layers are all empty from a particular value of k forward because of the triangular inequalities involving both bulk and boundary spins. It is the case of the Δ_3 , and the Δ_4 triangulations studied numerically in [11,23].

we could map the layer in an interval of integers, define a uniform distribution there and map it back to \mathcal{J}_k . This prescription is very unpractical. We prefer to define the uniform distribution in an alternative way.

We extract one real number from a continuous uniform distribution in $[0, k + 0.5]$ for each bulk face. We floor them to half-integers and we check if they belong to the layer \mathcal{J}_k . If they do, we accept them as a random sample of the layer. If they do not, we discard them and repeat the procedure. We summarize this procedure in the flowchart 1.

Algorithm 1. Random sampling of spins in \mathcal{J}_k .

```

1: procedure RANDOMSAMPLE( $k, \mathcal{J}_k$ )
2:   while true do  $\triangleright$  repeat until we find a good candidate
3:      $j_f \leftarrow$  extract a real number from a uniform distribution
       in  $[0, k + 0.5]$  for each face
4:      $j_f \leftarrow$  floor them to half-integers and interpret them as
       spins
5:     if all spins  $j_f$  are smaller than  $k$  then
6:       continue
7:     if any spin  $j_f$  do not satisfy triangular inequalities of  $\mathcal{J}_k$ 
       then
8:       continue
9:     return  $j_f$   $\triangleright$  the spins belong to the layer  $\mathcal{J}_k$ 

```

In the application we present in Sec. V, we perform an explicit test to show that the samples extracted with this algorithm are uniformly distributed in \mathcal{J}_k .

We acknowledge that Algorithm 1 is not optimal. To scale it up to more complicated amplitudes, we must improve it considerably. We leave this task to future work. Since they share the edge structure, the sampling Algorithm 1 is the same for both spin foam models we study.

We conclude this section by showing a pseudocode representation (Algorithm 2) of the Julia scripts that implement the Monte Carlo estimate of the spin foam amplitude (9). The full Julia scripts are available in the repository [30].

Algorithm 2. Scheme of Monte Carlo estimate of the spin foam amplitude.

```

1: for  $k = 0.5, 1 \dots K$  do  $\triangleright$  for each layer
2:   extract  $N_{mc}$  samples of the
   layer using RANDOMSAMPLE( $k, \mathcal{J}_k$ )
3:   store them in memory
4:
5: for  $k = 0.5, 1 \dots K$  do  $\triangleright$  for each layer
6:   load the bulk spins samples from memory
7:   for  $n = 1 \dots N_{mc}$  do  $\triangleright$  for each sample
8:     compute the partial amplitude  $a(j_f)$ 
9:     sum the partial amplitudes and save
       the layer amplitude  $S_k$ 
10:  compute the amplitude  $A_\Delta^{mc}(k)$ 
       summing the layer  $k$  to the previous ones
11:  store the amplitude

```

In calculating an EPRL spin foam amplitude, the partial amplitude also depends on the truncation parameter Δl . We fix the truncation parameter once and for all and store the amplitude for every Δl .

IV. APPLICATIONS TO THE MELONIC SELF-ENERGY AND VERTEX RENORMALIZATION DIAGRAMS

We test the effectiveness of the Monte Carlo framework described in Sec. III by computing four spin foam amplitudes. We focus on the melonic self-energy amplitude and the vertex renormalization (or 5–1 Pachner move) amplitude with the topological BF $SU(2)$ and the EPRL model. These diagrams are essential for studying

the infrared divergences of spin foam theories and their continuum limit.

The melonic self-energy diagram comprises two vertices, two boundary edges (one for each vertex), and four bulk ones connecting the two vertices, four boundary faces, and six bulk faces. We report in Fig. 1 a schematic representation of the 2-complex, and we refer to Appendix D for the detailed routing diagram.

To simplify the numerical calculation, we consider very symmetric boundary data where the boundary spins j_b are all the same and all the boundary intertwiners i_b are all the same. The spin foam transition amplitude with a homogeneous cutoff on all the spins associated with bulk faces is

$$A_{se}(j_b, i_b; K) = \sum_{j_f=0}^K \left(\prod_{f=1}^6 (2j_f + 1) \right) \sum_{i_e} \left(\prod_{e=1}^4 (2i_e + 1) \right) A_v(j_b, j_b, j_b, j_b, j_1, j_2, j_3, j_4, j_5, j_6; i_b, i_1, i_2, i_3, i_4) A_v(j_b, j_b, j_b, j_b, j_1, j_2, j_3, j_4, j_5, j_6; i_b, i_4, i_3, i_2, i_1). \quad (10)$$

The amplitude (10) is the same for both spin foam models that differ by the vertex amplitudes A_v : (2) for BF $SU(2)$ and (3) for EPRL. The melonic self-energy diagram with the BF $SU(2)$ topological model can be evaluated analytically and numerically [24,29,34]. We know it is divergent. The divergence is due to redundant $SU(2)$ delta functions that indicate some residual gauge freedom in the path integral [35] and can be dealt with by gauge fixing appropriately. The same amplitude with the EPRL model has been studied analytically [36], with a hybrid calculation [34], and, recently, numerically [24,29]. The amplitude is divergent, and there are strong indications that it diverges linearly in the cutoff. We are revisiting this amplitude as a control for the Monte Carlo technique we introduce. In fact, the self-energy has a relatively small number of internal

faces, so the computation is still possible even without using Monte Carlo. Reproducing known results allows us to evaluate the choices of the framework.

The vertex renormalization diagram contains five vertices, five boundary edges (one for each vertex), and ten bulk ones connecting all couples of vertices, ten boundary faces, and ten bulk faces. We report in Fig. 1 a schematic representation of the 2-complex, and we refer to Appendix D for the detailed routing diagram.

Also in this diagram, we simplify the numerical calculation by taking symmetric boundary data with all equal boundary spins j_b and boundary intertwiners i_b . If we put a homogeneous cutoff K on the sums over the bulk spins, the amplitude reads

$$A_{vr}(j_b, i_b; K) = \sum_{j_f=0}^K \left(\prod_{f=1}^{10} (2j_f + 1) \right) \sum_{i_e} \left(\prod_{e=1}^{15} (2i_e + 1) \right) A_v(j_b, j_b, j_b, j_b, j_1, j_2, j_3, j_4, j_5, j_6; i_b, i_4, i_{11}, i_{12}, i_2) A_v(j_b, j_b, j_b, j_b, j_7, j_8, j_1, j_9, j_2, j_3; i_b, i_6, i_{13}, i_{14}, i_4) A_v(j_b, j_b, j_b, j_b, j_{10}, j_4, j_7, j_5, j_8, j_1; i_b, i_8, i_{15}, i_{11}, i_6) A_v(j_b, j_b, j_b, j_b, j_6, j_9, j_{10}, j_2, j_4, j_7; i_b, i_{10}, i_{12}, i_{14}, i_{15}) A_v(j_b, j_b, j_b, j_b, j_3, j_5, j_6, j_8, j_9, j_{10}; i_b, i_2, i_{13}, i_{15}, i_{10}) \{6j\}(j_b, j_3, i_1, j_5, j_6, i_2) \{6j\}(j_b, j_1, i_3, j_2, j_3, i_4) \{6j\}(j_b, j_7, i_5, j_8, j_1, i_6) \{6j\}(j_b, j_{10}, i_7, j_4, j_7, i_8) \{6j\}(j_b, j_6, i_8, j_9, j_{10}, i_{10}) (-1)^\chi, \quad (11)$$




FIG. 1. Schematic representation of the 2-complexes of the self-energy spin foam diagram (left) and the vertex renormalization spin foam diagram (right).

where $\{6j\}$ are $SU(2)$ invariants that we define in Appendix A. The phase in (11) reduces to:

$$\chi = \sum_{k=1}^{10} j_k + \sum_{k=11}^{15} i_k. \quad (12)$$

The form of the amplitude is convoluted because we want to use the same intertwiner recoupling scheme in all the vertices. This is necessary to perform the numerical calculations efficiently, as `sl2cfoam-next` implements only a specific vertex amplitude (A4). Again, the amplitude (11) is the same for both spin foam models that differ by the vertex amplitudes. The vertex renormalization diagram can be evaluated analytically with the BF $SU(2)$ topological model integrating explicitly the group functions in the holonomy representation of the amplitude [34]. This amplitude was already studied with the Euclidean EPRL model in [37], finding a logarithmic divergence. A numerical calculation of the amplitude for values of the cutoff greater than 4 is extremely challenging if we do not use Monte Carlo. The degree of divergence of the same diagram with the EPRL spin foam model is entirely unknown. Any calculation with known techniques is too complicated. In Sec. VIII we study it using Monte Carlo. Computing this amplitude is a stress test for the Monte Carlo framework

and a novel result for studying EPRL spin foam infrared divergences.

In Fig. 2, we show the number of bulk spin configurations j_f as a function of the cutoff K for the vertex renormalization diagram. We only consider spin configurations that satisfy triangular inequalities. It is evident that the number of configurations to be summed increases as a power law with the cutoff K . Hence the convenience of using Monte Carlo. A simple numerical fit for $K \in [5, 10]$ shows that the number of configurations qualitatively scales as $\sim 39.3 \cdot K^{8.5}$ for the vertex renormalization. For the self-energy diagram, in [29], the same fit for $K \in [5, 20]$ showed that the number of configurations scales as $\sim 17.1 \cdot K^{5.6}$.

In all the amplitudes, to perform the numerical calculations with a modest amount of resources, we restrict the numerical calculation to the simplest nontrivial case of $j_b = \frac{1}{2}$ for both instances of boundary intertwiners $i_b = 0$ and $i_b = 1$. In the following sections, we will explicitly discuss the calculation's result only in the case of boundary intertwiner $i_b = 0$. However, we performed the same analyses also with boundary intertwiners $i_b = 1$. We find qualitatively identical results. Interested readers can find them in the detailed notebook in our public repository [30].

In the case of the EPRL model, we need to specify two more parameters to perform a numerical calculation. We fix the Immirzi parameter to $\gamma = 0.1$. We choose this value to partially compare our results with the literature on the numerical evaluation of EPRL spin foam amplitudes [24]. For similar reasons, we also choose the truncation parameter $\Delta l = 10$. This choice is also motivated by keeping the numerical task practical. The cost of resources increases rapidly with Δl , and literature [20,23,24] suggests that for maximal spins of order 10, the truncation $\Delta l = 10$ is a good compromise between costs and precision.

In the following sections, we use the same name for the amplitudes A_{se} and A_{vr} with both the BF $SU(2)$ and the EPRL model to keep the notation as clean as possible. The reader can uniquely identify which model the




FIG. 2. Number of bulk spins configurations j_f as a function of the cutoff K of the vertex renormalization diagram.

TABLE I. We compare the first three momenta of the sample with the corresponding exact quantities of a discrete uniform probability distribution. We find excellent agreement.

Quantity	Expected	Sample	Difference (%)
Mean	274 703.5	274 864.7	0.05%
Standard deviation	1 587 371.04	1 585 998.51	0.09%
Skewness	0	-0.00545	...

amplitude is computed with from the section. Finally, we use the term “exact amplitude” referring to (8) with a finite cutoff K computed without resorting to Monte Carlo methods.

V. THE MELONIC AMPLITUDE IN THE TOPOLOGICAL THEORY

In this section, we use the Monte Carlo framework described in Sec. III on the melonic self-energy transition amplitude in the topological $SU(2)$ BF model. This calculation aims to fine-tune and validate our choices of Monte Carlo parameters.

First, we test if the algorithm we use to sample the layers is equivalent to a uniform discrete probability distribution over the amplitude layers. We list *all* the sets of bulk spins in the layer, and we map them in an interval of integers.

Each element of the list is associated with its positions (we choose the order of the list arbitrarily but only once). We produce many samples using the Algorithm 1. We compute the samples’ mean, variance, and skewness and check if they are compatible with the corresponding quantities of a discrete uniform probability distribution. We tested every layer of this amplitude and found excellent agreement. For brevity, we report the analysis with a sample of 100 000 configurations of the amplitude layer with $k = 10$ that contain $V_{\mathcal{J}_k} = 549406$ possible configurations. We report them in Table I and Fig. 3. The probability distribution generated with Algorithm 1 is equivalent to a uniform discrete probability distribution.

We estimate the value and error of each amplitude layer with Monte Carlo repeating it $T = 20$ times and computing the mean and standard deviation. We choose the number of trials after a simple test. We fix the size of the Monte Carlo sampling to $N_{mc} = 1000$ to efficiently iterate and improve the analysis. We compute the average over $T = 10, 20,$ and 50 trials. We repeat it 100 times to study the distribution of the average. The law of large numbers states that the distribution of the averages is normal with standard deviation given by the average standard deviation. We can visualize it using a box plot we report in Fig. 4. This qualitative analysis shows that the tails of the distribution with 10 trials are very long. The standard deviation with just 10 trials is not a reasonable estimate of the error of the




FIG. 3. Histogram of the sample with 100 000 configurations from the layer $k = 10$ with 50 bins. Each bin contains approximately 2000 elements.




FIG. 4. Box plot of the average of the Monte Carlo evaluation of different amplitude layers. We consider the layers from $k = 5$ to 10. We repeat the estimate 100 times. We look at 10 trials (in blue), 20 trials (in orange), and 50 trials (in green). To ease the comparison, we plot the estimated value of the layer relative to the exact one.

Monte Carlo estimate of the amplitude's layer. This observation is independent of the layer. With 20 trials, the first and third quartiles are reduced to half, making it a better option. A similar observation is valid for 50 trials, albeit more resource intensive. We use 20 trials as a good compromise between precision and simplicity. Nevertheless, the error between the Monte Carlo estimate relative to the exact quantity is always a few percent with $N_{mc} = 1000$ and using 20 trials instead of 10 improve its estimate from 1% to 2%. The gain in the error estimate using 50 trials is marginal and does not justify the requirement of extra resources.

We study the Monte Carlo estimate of the amplitude as a function of the cutoff K for three different sample size choices of $N_{mc} = 1000$, $N_{mc} = 10000$, and $N_{mc} = 100000$. We average the calculation of $T = 20$ trials for each layer and sum them to get the amplitude. We compute the error on the amplitude from the standard deviation of each layer. We compare the relative error on the amplitude as a function of the cutoff for different sizes of Monte Carlo samples (see Fig. 5). For all three sample sizes, the relative error on the amplitude is smaller than 1%. As expected, the error decreases for larger values of N_{mc} . The relative error for $N_{mc} = 100000$ is smaller than 0.1%. We decide to use $N_{mc} = 100000$ for all the other calculations we present in this section. We also plot the estimated

value of the amplitude with $N_{mc} = 100000$ relative to the exact value with the estimated errors. We plot in Fig. 6 the Monte Carlo estimate of the amplitude in relation to the exact value, with error bars corresponding to the standard deviation. The exact value of the amplitude is compatible with the Monte Carlo estimate within the errors. At first sight, one could be confused by the trend of the errors decreasing with the cutoff. The observation that the relative error of the various layers is almost constant (as we can infer from Fig. 4) can easily explain this. However, the contribution to the amplitude of the outer layers (with larger k) is bigger than the others. A quick back-of-the-envelope calculation shows that if we add two quantities a few orders of magnitude apart but with the same relative error, the relative error on the sum is smaller than both.

We conclude our exploration by estimating the degree of divergence of this amplitude. The analytic calculation shows that the amplitude diverges with the cutoff as $\propto K^9$ at the leading order. Can we determine it numerically? We answer with a proof of concept analysis we use to validate the technique before applying it to more complex amplitudes where the analytic answer is unknown. We are not satisfied with a qualitative result. We could easily eyeball a line on the logarithmic plot of the amplitude as a function of the cutoff. However, this approach is only helpful if we know the degree of divergence. We need to




FIG. 5. Relative error of the $SU(2)$ BF melonic amplitude as a function of the cutoff computed with 20 trials. We compare different Monte Carlo sampling sizes $N_{mc} = 1000$ (green), $N_{mc} = 10000$ (orange), and $N_{mc} = 100000$ (blue) samples.




FIG. 6. Monte Carlo estimates of the $SU(2)$ BF melonic amplitude relative to the exact value with $N_{mc} = 100000$ and $T = 20$ trials in each layer.




FIG. 7. Monte Carlo estimates of the $SU(2)$ BF melonic amplitude with $N_{mc} = 100\,000$ and $T = 20$ trials in each layer. We compare different values of the parameter μ in the face amplitude (14).

perform a fit to determine it numerically. We start with a model function with a polynomial form

$$A(K) = c_1 K^a + c_2 K^{a-1}. \quad (13)$$

We limit ourselves to the leading and subleading order terms. In general, the amplitude diverges as a polynomial with all the powers of the cutoff. Using it as a model to fit our data would undoubtedly lead to overfitting as we want to use a maximum cutoff of $K = 10$. We could compute this amplitude for larger values of the cutoff. However, the EPRL model's amplitude is too computationally demanding, and we must impose a small cutoff of $K = 10$. We use this limitation as an excuse to use a small cutoff with the topological model and anticipate some problems arising from this choice. We fit using only the last 10 data points available.

We perform a simple least squares fit using the Julia package `LsqFit`. We find the exponent $a = 8.81 \pm 11.99$ and coefficients $c_1 = 4.83 \pm 240.38$, $c_2 = 41.12 \pm 66.95$. Examining the uncertainties of the parameters, we conclude that the fit is clearly unreliable. Moreover, even if the fit value for the exponent a looks compatible with the exact value $a = 9$, we could not affirm it without knowing it in advance.

What is happening? The covariance between the coefficients c_1 and c_2 is huge. The fit procedure with a cutoff of order 10 cannot distinguish between the contributions from the leading and subleading orders (for example, if the exact coefficients are $c_1/c_2 \approx 10$).

A simple solution to this impasse would be to extend the fit to larger values of K . In this way, the contribution of the leading order would dominate the subleading one. Since we

are limited by a maximum cutoff $K = 10$, we have to find a different solution. Alternatively, we can diminish the degree of divergence of the amplitude by changing the face amplitude introducing a tunable parameter μ

$$A_f(j_f) = (2j_f + 1) \rightarrow (2j_f + 1)^\mu. \quad (14)$$

The case $\mu = 1$ corresponds to the standard case, but if we set $\mu < 1$ we lower the divergence of the amplitude. In general, the amplitude will diverge as $A_{se} \propto K^{6\mu+p}$ where p is a number we have to determine, and 6μ is the contribution coming from the six unbounded sums over the bulk spins. We pretend we do not know that for $\mu = 1$ the amplitude diverges as $A_{se} \propto K^9$ and therefore $p = 3$. And we try to determine $p = a - 6\mu$ fitting the amplitude with the same model (13) for various values of $\mu = 1/6$, $\mu = 0$ and $\mu = -1/6$. We take this opportunity to check if the Monte Carlo estimate of the amplitude is as good as in the case $\mu = 1$. We compare the Monte Carlo estimate of the amplitudes relative to their exact values with different μ . All the relative error bars, computed over $T = 20$ realizations as we did before, are within 0.1%, confirming that the Monte Carlo estimate is very accurate. We summarize the results in Fig. 7. We fit the amplitude with the model (13). The interpretation of the result, in this case, is more straightforward. We report the fitted coefficients in Table II.

A few comments are in order. First, all the fits indicate clearly that $p = a - 6\mu = 3$. Second, notice that we are not worried that a is not always compatible with the nearest integer value. This is an artifact of using just the leading order and next to the leading order of the polynomial in (13). We determined the degree of divergence of the amplitude as $A_{se} \propto K^{6\mu+3}$.

VI. THE VERTEX RENORMALIZATION AMPLITUDE IN THE TOPOLOGICAL THEORY

Exact numerical calculations of spin foam amplitudes with many bulk faces are accessible only for simple models, but become infeasible when the number of faces is too large. We showcase the problem by looking at the vertex renormalization or 5-1 Pachner move amplitude with

TABLE II. Values of the coefficients of the model (13) obtained fitting the amplitude with different values of μ .

	a	c_1	c_2
$\mu = 1/6$	4.15 ± 0.34	0.91 ± 1.24	6.38 ± 0.17
$\mu = 0$	2.97 ± 0.04	2.01 ± 0.03	3.50 ± 0.05
$\mu = -1/6$	1.94 ± 0.02	2.83 ± 0.02	0.39 ± 0.05




FIG. 8. Relative error on the Monte Carlo estimate of the $SU(2)$ BF vertex renormalization amplitude over 20 trials. We compare different Monte Carlo sampling sizes $N_{mc} = 1000$ (green), $N_{mc} = 10\,000$ (orange), and $N_{mc} = 100\,000$ (blue) samples.

the topological $SU(2)$ BF spin foam theory. The issue with this computation is not the time we need to compute each term of the sums over the spins of the ten internal faces but the sheer amount of terms of these sums. With a cutoff $K = 10$ on the sums, we need to compute 307 88 382 715 terms in total, 11 892 969 195 of which belongs to last layer of the amplitude. They are almost six orders of magnitude more than in the melonic diagram case. This is where using the Monte Carlo framework to perform the sums is necessary.

Motivated by the analysis of the melonic diagram, we average the Monte Carlo calculation of each layer over $T = 20$ trials. We sum the layers' average to obtain the amplitude value for a given cutoff. We compute the amplitude variance by summing each layer's variance. We consider the standard deviation as the error of the amplitude. We perform the calculation with three different choices of sample sizes $N_{mc} = 1000$, $N_{mc} = 10\,000$, and $N_{mc} = 100\,000$.

In this case, we cannot compare with the exact value of the amplitude to evaluate the Monte Carlo technique. The exact value is not computable for cutoff $K = 10$. It is exactly the reason we resort to Monte Carlo methods. We compare the relative error on the amplitude for the three sample size choices. We summarize the result of this analysis in the plot of Fig. 8. We see that $N_{mc} = 1000$ produces an estimate with a significant relative error between 8% and 3%. For $N_{mc} = 100\,000$ we find very modest relative errors between 0.8% and 0.3%. The increase in the relative errors compared to the melonic diagram case is expected. Each layer of this amplitude contains a few orders of magnitude more elements than the corresponding layer in the melonic amplitude.

TABLE III. Values of the coefficients of the model (13) obtained fitting the amplitude A_{vr} with different values of μ .

	a	c_1	c_2
$\mu = 0$	1.96 ± 0.02	0.18 ± 0.01	-0.05 ± 0.02
$\mu = 1/10$	2.89 ± 0.02	0.15 ± 0.01	0.07 ± 0.02
$\mu = 1/5$	4.00 ± 0.04	0.08 ± 0.01	0.26 ± 0.02

We fit the amplitude computed with $N_{mc} = 100\,000$ samples averaged over $T = 20$ trials as a function of the cutoff K using the model (13). We employ only the amplitude value as a function of the cutoff between 5 and 10 since we expect the matching of the functional form (13) to be, at best asymptotic. The degree of divergence of this amplitude can be estimated analytically, finding $A_{vr} \propto K^{12}$ at the leading order.

With a simple least squares fit, we find an unreliable result with an exponent $a = 11.20 \pm 0.69$ and coefficients $c_1 = 1.91 \pm 4.39$, and $c_2 = 1.70 \pm 10.82$. The situation is analogous to the case of the melon amplitude. To determine numerically the degree of divergence of this amplitude with a maximum cutoff of $K = 10$ we modify the face amplitude as in (14) such that the amplitude diverge as $A_{vr} \propto K^{10\mu+p}$ with p to determine. We want to keep the exponent $10\mu + p$ as low as possible. Therefore, we pick three values of the weight $\mu = 0$, $\mu = 1/10$, and $\mu = 1/5$. We find all three results are compatible with the analytic value of $p = 2$ resulting in the amplitude diverging as $A_{vr} \propto K^{10\mu+2}$ (see Table III).

VII. THE MELONIC AMPLITUDE IN THE EPRL THEORY

In Sec. V, we computed the melonic self-energy spin foam amplitude with the topological BF $SU(2)$ model using Monte Carlo. We obtained a remarkably accurate amplitude approximation, employing only a fraction of the computational resources. Is the Monte Carlo technique applicable to spin foam amplitudes with the EPRL model, and is it equally successful?

Before discussing the calculation details, we must disentangle two different overlapping approximations. One is due to the Monte Carlo sampling procedure, while the other is a consequence of working with a finite truncation parameter Δl . In calculating the melonic amplitude, we have access to the public data from [24] that employs an extensive truncation $\Delta l = 20$. We borrow that data to perform a detailed study of the truncation approximation independently from the Monte Carlo one.

We mitigate the dependence from a specific choice of truncation Δl using an extrapolation technique (see Appendix C). This idea was first introduced in [24], further formalized in [23], and additionally utilized in [25]. Our understanding of the extrapolation technique applied to EPRL spin foam amplitudes with finite truncation has improved considerably. In this section, we revisit it in a new

light. We explicitly show that its principal hypothesis is satisfied and compare possible alternatives.

The amplitude $A_{se}(K, \Delta l)$ at fixed cutoff K is a sequence in the truncation parameter Δl . Since the EPRL vertex amplitude is well-defined, the limit of infinite truncation is finite, and we can approximate it using the Aitken delta squared method.

$$A_{se}^{(ex)}(K, \Delta l) = \frac{A_{se}(K, \Delta l)A_{se}(K, \Delta l - 2) - A_{se}(K, \Delta l - 1)^2}{A_{se}(K, \Delta l) - 2A_{se}(K, \Delta l - 1) + A_{se}(K, \Delta l - 2)}. \quad (15)$$

The sequence $A_{se}^{(ex)}(K, \Delta l)$ converges to $A_{se}(K)$ faster than $A_{se}(K, \Delta l)$. Therefore approximating the limit with the truncation of the sequence $A_{se}(K) \approx A_{se}^{(ex)}(K) \equiv A_{se}^{(ex)}(K, \Delta l)$ is, in general, a better approximation than using the truncation of $A_{se}(K, \Delta l)$. The extrapolation is effective if the rate of convergence of the amplitude is at least linear in Δl , as discussed in Appendix C. For this purpose, we study the ratio

$$\lambda_{A_{se}(K)}(\Delta l) = \frac{A_{se}(K, \Delta l) - A_{se}(K, \Delta l - 1)}{A_{se}(K, \Delta l - 1) - A_{se}(K, \Delta l - 2)}. \quad (16)$$

If the limit of the ratio (16) for infinite truncation is smaller than 1, the convergence of $A_{se}(K, \Delta l)$ is linear. Proving numerically the existence of the limit is challenging. Therefore, we settle with some numerical evidence for linear convergence. Moreover, truncating the sequence of extrapolations $A_{se}^{(ex)}(K, \Delta l)$ to a finite Δl to approximate its limit is reliable only if the ratio approached (at least approximately) a horizontal asymptote. Part of this analysis has already been performed in [24]. Here we repropose it in light of our improved understanding. For all half-integers $K \leq 10$, the ratio (16) approaches a horizontal asymptote smaller than 1. This behavior is evident for any $8 \leq \Delta l \leq 20$ and legitimizes the extrapolation of the amplitude (15). The ratio approaches the asymptote from below. Therefore, we expect the extrapolations from larger

truncation to increase. We summarize the analysis for some cutoff values in Fig. 9.

How effective is the extrapolation (15)? We answer this question by comparing the extrapolation obtained from $\Delta l = 10$ and $\Delta l = 20$. The last is usually inaccessible due to its exceptional computational cost, while the former is less precise but cheaper to compute. The value of the amplitude $A_{se}(K, \Delta l)$ changes substantially when we increase the truncation. For example, at cutoff $K = 10$, the amplitude with $\Delta l = 20$ is 40% larger than the one with $\Delta l = 10$. This was expected since the spins entering the calculation at $K = 10$ are of order 20, and the amplitude truncated at $\Delta l = 10$ cannot approximate the real amplitude value well. However, we find a milder difference between the extrapolation (15) with different truncations. With the same cutoff, the difference between the extrapolation done with truncation $\Delta l = 10$ and $\Delta l = 20$ is just 6%. In both cases, the extrapolations approximate the amplitude better than any truncated amplitude we have access to. We summarize the results in Fig. 10.

It is natural to question if extrapolating the amplitude sequence (15) is the only way to proceed. We explored multiple alternatives (vertex-by-vertex, bulk amplitude, and layer-by-layer), and they all turned out to perform worst. We review in detail one of them. The amplitude is a sum limited to a maximum K of contributions of layers (8) we compute with a fixed truncation Δl




FIG. 9. Plot of the ratio (16) for various values of the cutoff $K = 10$, $K = 7.5$, $K = 5$. A horizontal asymptote is evident even for modest values of the truncation.




FIG. 10. Comparison between the amplitude extrapolations and the truncated values $A_{se}(K, \Delta l)$ with $\Delta l = 10$ and $\Delta l = 20$ as a function of the cutoff. For larger cutoff values, the effect of truncation is critical. The extrapolation successfully mitigates this dependence.

$$A_{se}(K, \Delta l) = \sum_{k=0}^K S_k(\Delta l). \quad (17)$$

Each layer contribution $S_k(\Delta l)$ is a convergent sequence in the truncation parameter. We can use Aitken extrapolation on each layer contribution and obtain a sequence of amplitudes summing them

$$A_{se}^{(ex,S)}(K, \Delta l) = \sum_{k=0}^K S_k^{(ex)}(\Delta l) \quad \text{with} \quad (18)$$

$$S_k^{(ex)}(\Delta l) = \frac{S_k(\Delta l)S_k(\Delta l-2) - S_k(\Delta l-1)^2}{S_k(\Delta l) - 2S_k(\Delta l-1) + S_k(\Delta l-2)}.$$

We approximate the limit of the sequence of amplitude truncating the accelerated convergence sequence $A_{se}^{(ex,S)}(K) \approx A_{se}^{(ex,S)}(K, \Delta l)$ (18). Since the amplitude (17) is a finite sum over layers, the sequence obtained extrapolating layer-by-layer (18) has the same limit as the sequence (15) and both, of course, converge to the value of the amplitude $A_{se}(K)$ without any truncation. In practice, we do not have access to an arbitrarily large truncation but to a relatively small one, and we want to approximate the limit truncating the extrapolated sequences. Which one approximates the amplitude better is an

open question that we can answer by exploring different possibilities.

First, we verify that every layer amplitude is at least linearly convergent in the truncation parameter by studying the ratio

$$\lambda_{S_k}(\Delta l) = \frac{S_k(\Delta l) - S_k(\Delta l-1)}{S_k(\Delta l-1) - S_k(\Delta l-2)}. \quad (19)$$

We show some explicit examples in Fig. 11. All the layer amplitudes are compatible with the linear convergence hypothesis. Differently from (16) the ratio (19) approach a horizontal asymptote from above. As a consequence, the extrapolation of the layer amplitudes decreases in value for increasing truncation. In particular, the last layer $k = 10$ for $\Delta l = 10$ is still far from the horizontal asymptote. For this reason, we expect the extrapolation to change significantly if we increase the truncation from $\Delta l = 10$ to, for example, $\Delta l = 15$ or 20 and approximate. This is not the case for lower levels $k \leq 5$. We extrapolate all the layer amplitudes and sum them. We truncate the sequence (18) to approximate its limit. It is useful to compare the result with the extrapolation of the whole amplitude (15) using different truncation parameters. We summarize our findings in Fig. 12. For large cutoff, the extrapolation (18) with




FIG. 11. Plot of the ratio (19) for various values $k = 10$, $k = 7.5$, $k = 5$. A horizontal asymptote is evident but is reached for different truncation values.




FIG. 12. Comparison of the melonic self-energy amplitude obtained extrapolating the amplitude layer-by-layer or the whole amplitude.

truncation $\Delta l = 10$ is much larger than its counterpart (15). At cutoff $K = 10$, the layer-by-layer extrapolation (18) is 16% larger than whole amplitude extrapolation (15). The difference reduces drastically if we repeat both extrapolations with a larger truncation $\Delta l = 20$. At cutoff $K = 10$ the two differ only by 1%. In particular, the value of the extrapolation (18) decreases substantially. This agrees with what we observed studying the ratio (19). With low truncation, the layer-by-layer extrapolation results in a poor amplitude approximation. In contrast, the two extrapolation schemes almost coincide for larger truncations.

In the following, we will study the Monte Carlo approximation of the amplitude and limit ourselves to a small truncation $\Delta l = 10$ to conserve computational resources. We will use only the whole amplitude extrapolation scheme (15) since it is the most accurate within this setting.

Next, we evaluate how effective Monte Carlo techniques are if applied to the bulk spin summations in the EPRL model using the melonic amplitude as a testing ground. Following the cost-benefit analysis of the first part of this section, we set the truncation to $\Delta l = 10$. We use Monte Carlo to estimate the contribution to the amplitude of each layer averaging over $T = 20$ trials. The amplitude is given by the sum of the averages of the layers and its error by the square root of the total variance.

We perform the calculation with three different choices of Monte Carlo sample sizes $N_{mc} = 1000$, $N_{mc} = 10\,000$, and $N_{mc} = 100\,000$. As displayed in Fig. 13 the relative error on the amplitude is more or less stable at 1% for the small sample size and 0.1% for the large one. The error is compatible with the analog error computed with the topological theory. We could have expected it as the distribution of the layers' amplitudes value is relatively flat. The dominant factor in the error is the ratio between the number of configurations in the layer and the Monte Carlo sample size, which is model-independent. At fixed truncation $\Delta l = 10$, we can also compare the amplitude computed using Monte Carlo with the exact one. The amplitude computed using Monte Carlo is compatible with the exact value within the error. We summarize the comparison in the plot in Fig. 14.

How well the extrapolation technique (15) is compatible with the Monte Carlo sum over the bulk spins? Instead of averaging over $T = 20$ trials, we sum the layer amplitudes of each realization to obtain 20 different realization of the amplitude at fixed truncation $\Delta l = 10$. This is possible since each Monte Carlo estimate of each layer amplitude is independent. We extrapolate the amplitude (15) for each trial and approximate it averaging over the trials and considering as error its standard deviation.




FIG. 13. Relative error of the EPRL melonic amplitude as a function of the cutoff computed with 20 trials and truncation fixed to $\Delta l = 10$. We compare different Monte Carlo sampling sizes $N_{mc} = 1000$ (green), $N_{mc} = 10\,000$ (orange), and $N_{mc} = 100\,000$ (blue) samples.




FIG. 14. Monte Carlo estimate of the EPRL melonic amplitude with finite truncation $\Delta l = 10$ relative to the exact value with $N_{mc} = 100\,000$ and $T = 20$ trials in each layer.




FIG. 15. Monte Carlo estimates of the EPRL melonic amplitude extrapolated from a finite truncation $\Delta l = 10$ relative to the exact value with $N_{mc} = 100\,000$ and $T = 20$ trials in each layer.

We compare the extrapolation of the amplitude computed with Monte Carlo with the one calculated without that approximation. We find a 0.1% average error due to Monte Carlo for all cutoff values (see Fig. 15). This is compatible with the finite truncation case with $\Delta l = 10$. However, this is just an error due to the Monte Carlo approximation. We expect it to be firmly subdominant with respect to the error due to the presence of the truncation despite the extrapolation.

Finally, we determine the degree of divergence of the amplitude performing a numerical fit. The same calculation was already done in [24] with an exact amplitude computation. We showed that using Monte Carlo allows us to estimate the amplitude with a 0.1% error compared to the exact value (at fixed truncation). Therefore is no surprise that we find a good fit with

$$A(K) = c_1 K^a + c_2 \quad (20)$$

with $a = 1.091 \pm 0.005$, $c_1 = (6.186 \pm 0.093) \times 10^{-6}$, and $c_2 = (-3.476 \pm 0.253) \times 10^{-6}$ that coincide with the result of [24]. We changed the model of the fit from the topological models, as (13) gives unreliable results. If we generalize the face amplitude (14) introducing a weight μ , we conclude that the scaling of the amplitude is compatible with $A(K) \propto K^{6\mu+p}$ with $p \approx -5$.

For completeness, we could look at the melonic divergence with different weight μ values as we did for the

topological $SU(2)$ model. However, the result we obtain is unreliable. A more solid calculation requires a larger truncation and, consequently, way more computational resources that we currently do not have access to. We will comment on these issues in more detail in the next section. We leave this interesting consistency check to future work.

VIII. THE VERTEX RENORMALIZATION AMPLITUDE IN THE EPRL THEORY

Computing the degree of divergence of the EPRL vertex renormalization amplitude (11) has never been attempted. The theory is too complex to do it numerically or analytically. We can use Monte Carlo to calculate this amplitude for the first time. We build upon the experience accumulated in the previous sections and our choices and approximations. For this reason, we use a truncation parameter $\Delta l = 10$, and $N_{mc} = 100\,000$ Monte Carlo samples. The calculation of this amplitude required ~ 400 CPU hours, which is a minimal fraction of what would be required without Monte Carlo. We compute the Monte Carlo error as the standard deviation of the amplitude over $T = 20$ relative to its average. We find a very stable error of approximately 0.9% for all values of the cutoff as summarized in Fig. 16. The relative error of the amplitude is curiously constant for increasing cutoff K . While this behavior could seem odd initially, it has a straightforward explanation. As we discuss below, the amplitude seems convergent. The amplitude layers




FIG. 16. Relative error of the EPRL vertex renormalization amplitude as a function of the cutoff, computed with 20 trials and truncation fixed to $\Delta l = 10$. We use $N_{mc} = 100\,000$ Monte Carlo samples.




FIG. 17. Monte Carlo estimate of the EPRL vertex renormalization amplitude as a function of the cutoff K . We use $\Delta l = 10$, $N_{mc} = 100\,000$, and $T = 20$ trials. We plot the extrapolated amplitude and the fit using the model (21).

decrease very fast. The contribution to the relative error of larger layers is irrelevant.

The amplitude's statistical fluctuations are slightly bigger than in the EPRL self-energy amplitude. However, they are compatible with the errors of the topological model. It is not surprising since the ratio between the number of samples and the cardinality of the set we are summing over dominates the error of a Monte Carlo calculation.

We reduce the dependence of the amplitude from the truncation using the extrapolation (15) as discussed in Sec. VII. We show the value of the extrapolated amplitude as a function of the cutoff in Fig. 17. The amplitude is essentially constant for cutoff $K > 2$. A power law fit is inadequate to capture the functional scaling of the amplitude. Therefore we opt for a model capturing the constant behavior plus a correction.

$$A(K) = c_1 + c_2/K. \quad (21)$$

Fitting the amplitude result in $c_1 = 0.765 \pm 2.667 \times 10^{-5}$ and $c_2 = -0.0006 \pm 0.0002$. We should take these values with a grain of salt as they depend strongly on the model we decide to use. We are tempted to enhance the divergence of the amplitude by modifying the face amplitude (14) introducing the weight μ . We observe that increasing μ we

need to increase Δl . Otherwise, the extrapolation technique fails in estimating the amplitude well.

We tried different face amplitudes weights and studied the amplitude ratios (16). We show it in Fig. 18. While for standard face amplitude $\mu = 1$ and $\mu = 1.4$, the ratio reaches a horizontal asymptote smaller than 1 very soon, it is not the case for $\mu = 1.8$, and $\mu = 2$. We can see how for these two cases, the ratio is still decreasing and greater than 1 at truncation $\Delta l = 10$. Therefore, to obtain a reliable extrapolation we would need a larger truncation, not accessible with the computational resources at our disposal.

We conclude that for this transition amplitude, the extrapolation scheme (15) is sensible to the weight on the face amplitudes. The slower convergence of the amplitude sequence is also evident from the plot of the amplitude for different truncations as well displayed in Fig. 19. For completeness, we report the fit of the convergent amplitude with $\mu = 1.4$ with the model (21). We find $c_1 = 5.535 \pm 0.011$ and $c_2 = -0.464 \pm 0.074$. Due to the invalidity of the extrapolation, we cannot perform a fit in the other two cases. Consequently, we cannot estimate p in the divergence of the amplitude $A_{vr} \propto K^{10\mu+p}$. Nevertheless, our numerical analysis shows strong indications that the infrared bubble of the EPRL theory is convergent.




FIG. 18. Plot of the ratio (16) for increasing values of μ at cutoff $K = 10$. For $\mu = 1.8$ and $\mu = 2$ the ratio has not reached a horizontal asymptote smaller than 1 at $\Delta l = 10$. Therefore the extrapolation (15) cannot be applied.

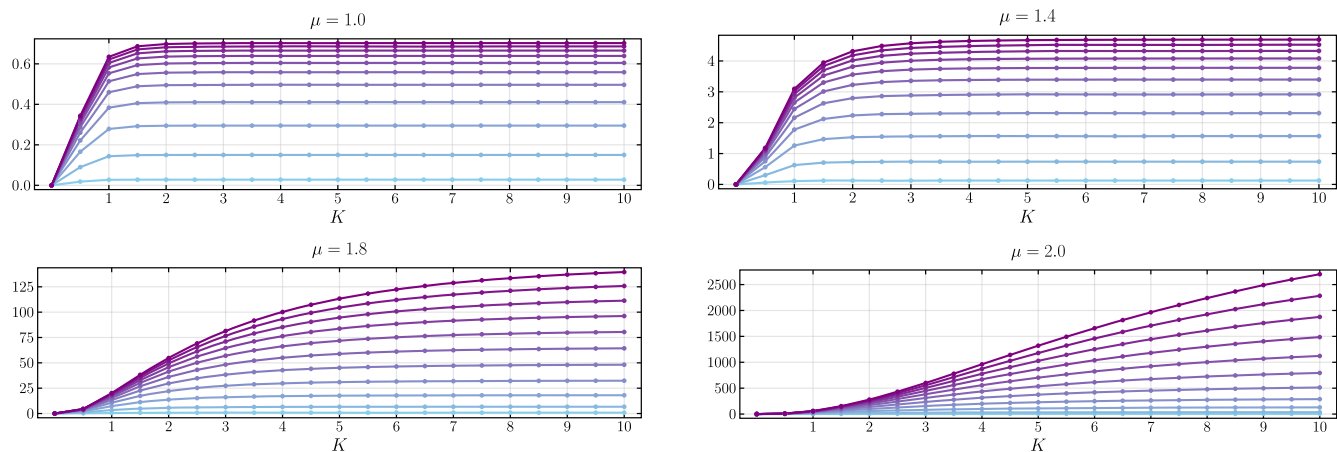


FIG. 19. Plots of vertex renormalization EPRL amplitude $A_{br}^{mc} \times 10^{17}$ for increasing values of the weight factor μ . In each panel, we report all the curves obtained for increasing values of the truncation parameter Δl . The bottom curve (azure) corresponds to $\Delta l = 0$ while $\Delta l = 10$ is the top one (purple).

IX. CONCLUSION AND DISCUSSION

The `sl2cfoam-next` library allows fast and reliable calculations of EPRL spin foam transition amplitudes. While it is optimized to compute vertex amplitudes, calculating a spin foam amplitude with many vertices and internal faces still presents a huge technical obstacle. To sum over the bulk degrees of freedom, we have to compute an enormous number of components that scale exponentially with the number of internal faces.

We apply Monte Carlo to the spin foam bulk summations and show that it is a very promising strategy to overcome this obstacle. The complexity of the calculation depends on the number of Monte Carlo samples N_{mc} we can freely choose. Of course, the result's precision depends on N_{mc} and how we choose the probability distribution of the Monte Carlo sampling. We decide to use a uniform probability distribution. We acknowledge it is not the optimal choice as it equally weights all the bulk spins configurations. However, it is efficient and allows us to parallelize the sampling algorithm for a single amplitude

across multiple threads. Moreover, we can use it democratically with any amplitude. Alternatively, we could abandon uniform sampling in favor of Markov Chain Monte Carlo (MCMC) methods. This would result in faster convergence, but on the other hand, the sampling algorithm for calculating a single amplitude would no longer be parallelizable. We leave the study and implementation of MCMC to bulk degrees of freedom for future works. Finally, in this paper, we used intertwiners as boundary states. In other cases, one could attempt to take advantage of the state's properties to perform importance sampling Monte Carlo. For example, in the case of extrinsic boundary states, it would make sense to tailor the probability distribution sampling from a normal distribution. One has to deal with the well-known sign problem with highly oscillatory distributions. We leave the analysis of different boundary states for future works.

We evaluate the proposed strategy and discuss its choices for computing the melonic self-energy and the vertex renormalization spin foam amplitudes, with the $SU(2)$

BF and the EPRL theory. These are amplitudes with many vertices and internal faces, providing a good test. In particular, the topological theory is convenient as partial analytical calculations are possible and help us evaluate the performance of our method. The calculations of the amplitudes with Monte Carlo are surprisingly effective already with a modest number of Monte Carlo samples $N_{mc} = 100\,000$. To appreciate this result remember that the vertex renormalization amplitude contains $\approx 10^{10}$ possible (nontrivially vanishing) amplitudes. Obtaining a good approximation computing just order $\approx 10^5$ of them is remarkable. We find small uncertainties of 0.1% and 0.9% for the EPRL self-energy and the vertex renormalization amplitude, respectively. Similar results hold for the $SU(2)$ BF model suggesting that the dominant factor in the error is the ratio between the cardinality of the space of all the possible spin configurations and the (square root of the) Monte Carlo samples. The calculations in the EPRL theory are carried out with a finite truncation $\Delta l = 10$. We alleviate the dependence of the result using extrapolation techniques to accelerate the convergence of the amplitude. We explore different extrapolation schemes. Extrapolating the full amplitude at a finite cutoff is the most convenient option. We also formalize the regime of validity of the extrapolation scheme and develop a test to verify if the amplitude falls into it.

At the same time, analyzing the divergence of these amplitudes for the EPRL theory is an essential step toward understanding the theory's continuum limit. While the self-energy amplitude is already studied in the literature [24,34,36], for the vertex renormalization amplitude, we know only a loose upper bound [34]. The amplitude was too complex to try any numerical or analytical calculation. The estimate we provide in this paper is a complete and important novelty. Performing the sums over the bulk degrees of freedom with statistical methods is enough to confirm the linear divergence of the melonic self-energy amplitude.

The numerical evaluation of the EPRL vertex renormalization spin foam amplitude provides a convincing argument to claim its convergence. It is a shocking result as it contradicts any intuition we could get from the analytical calculations of the topological models where the amplitude is more divergent than the melonic self-energy one.

Our result is a numerical computation of the amplitude, not analytical proof. One should always keep in mind its limitations. We performed a calculation with fixed boundary spins $j_b = \frac{1}{2}$, Immirzi parameter $\gamma = 0.1$, an extrapolation based on the truncation $\Delta l = 10$, and a uniform cutoff on all the faces limited to $K \leq 10$. Technical limitations and convenience dictate some choices. Nevertheless, we explored alternatives when possible,

and the result seems general. Numerically we infer that the convergence of the vertex renormalization amplitude is determined by the destructive interference of the vertex amplitudes' oscillations of the EPRL theory. The booster functions are responsible for the interference, which appear in the EPRL vertex amplitude and are not present in the topological one. They encode the imposition of the simplicity constraints in the theory. This interpretation agrees with the results in [34]. Neglecting this interference results in a divergent upper bound estimate identical to the topological model.

Consider the Ponzano Regge model, a simpler spin foam theory that describes euclidean quantum gravity in three dimensions. The vertex renormalization amplitude in that theory (the 1–4 Pachner move) is cubically divergent in the cutoff K^3 . The divergence is related to a residual gauge invariance in the path integral that is not entirely fixed [35]. Geometrically it can be interpreted in the ∞^3 ways we can divide a tetrahedron in four with an extra point. The convergence of the EPRL vertex renormalization amplitude could signal that a similar symmetry is not present in the theory. The restriction to Lorentzian geometries with spacelike boundaries of EPRL vertex in the large spins regime breaks the BF action's shift symmetry. Whether or not it indicates that simplicity constraints are imposed correctly in the EPRL theory remains an open question. We hope that a detailed analytical study and the contribution of other upcoming numerical techniques tailored to the study of the large spin limit of the theory [12,16] could help solve this mystery.

ACKNOWLEDGMENTS

P. D. is supported by the ID# 62312 grant from the John Templeton Foundation, as part of the “The Quantum Information Structure of Spacetime (QISS)” Project (qiss.fr). P. F. is supported by the Natural Science and Engineering Council of Canada (NSERC) through the Discovery Grant “Loop Quantum Gravity: from Computation to Phenomenology.” We also acknowledge the Shared Hierarchical Academic Research Computing Network (SHARCNET) and the Digital Research Alliance of Canada ([38]) for granting access to their high-performance computing resources. P. D. thanks Laurent Freidel for an insightful discussion on the consequence of the convergence of the EPRL vertex amplitude. P. F. thanks Francesco Gozzini for the useful discussions about this project. We acknowledge the Anishinaabek, Haudenosaunee, Lūnaapéewak and Attawandaron peoples, on whose traditional lands Western University is located.

APPENDIX A: VERTEX AMPLITUDES

In this appendix, we report the definition of the topological BF SU(2) spin foam vertex amplitude.

$$A_v^{BF}(j_1, j_2, j_3, j_4, j_5, j_6, j_7, j_8, j_9, j_{10}; i_1, i_2, i_3, i_4, i_5) = \begin{Bmatrix} i_1 & j_3 & i_4 & j_6 & i_2 \\ j_4 & j_{10} & j_8 & j_5 & j_1 \\ j_7 & i_5 & j_9 & i_3 & j_2 \end{Bmatrix}. \quad (\text{A1})$$

The SU(2) invariant in (A1) is a $\{15j\}$ symbol of the first kind. According to the conventions of [39], we write it in terms of Wigner's $\{6j\}$ symbols.

$$\begin{aligned} \begin{Bmatrix} j_1 & j_2 & j_3 & j_4 & j_5 \\ l_1 & l_2 & l_3 & l_4 & l_5 \\ k_1 & k_2 & k_3 & k_4 & k_5 \end{Bmatrix} &= (-1)^{\sum_{i=1}^5 j_i + l_i + k_i} \sum_x (2x+1) \begin{Bmatrix} j_1 & k_1 & x \\ k_2 & j_2 & l_1 \end{Bmatrix} \begin{Bmatrix} j_2 & k_2 & x \\ k_3 & j_3 & l_2 \end{Bmatrix} \\ &\times \begin{Bmatrix} j_3 & k_3 & x \\ k_4 & j_4 & l_3 \end{Bmatrix} \begin{Bmatrix} j_4 & k_4 & x \\ k_5 & j_5 & l_4 \end{Bmatrix} \begin{Bmatrix} j_5 & k_5 & x \\ j_1 & k_1 & l_5 \end{Bmatrix}. \end{aligned} \quad (\text{A2})$$

In this work we also use the in-line notation for

$$\{6j\}(j_1, j_2, j_3, j_4, j_5, j_6) = \begin{Bmatrix} j_1 & j_2 & j_3 \\ j_4 & j_5 & j_6 \end{Bmatrix}. \quad (\text{A3})$$

In the booster function decomposition, the Lorentzian EPRL vertex amplitude is defined as:

$$\begin{aligned} A_v(j_1, j_2, j_3, j_4, j_5, j_6, j_7, j_8, j_9, j_{10}; i_1, i_2, i_3, i_4, i_5) \\ = \sum_{l_f=j_f}^{\infty} \sum_{k_e} \begin{Bmatrix} i_1 & j_3 & k_4 & l_6 & k_2 \\ j_4 & l_{10} & l_8 & l_5 & j_1 \\ l_7 & k_5 & l_9 & k_3 & j_2 \end{Bmatrix} (2k_2+1)(2k_3+1)(2k_4+1)(2k_5+1) \\ B_4^{\gamma}(j_5, j_6, j_7, j_1, l_5, l_6, l_7, j_1; i_2, k_2) B_4^{\gamma}(j_8, j_9, j_2, j_5, l_8, l_9, j_2, l_5; i_3, k_3) \\ B_4^{\gamma}(j_{10}, j_3, j_6, j_8, l_{10}, j_3, l_6, l_8; i_4, k_4) B_4^{\gamma}(j_4, j_7, j_9, j_{10}, j_4, l_7, l_9, l_{10}; i_5, k_5). \end{aligned} \quad (\text{A4})$$

The booster functions are one-dimensional integrals over the rapidity parameter r of the reduced matrix elements, in the γ -simple unitary representation of $SL(2, \mathbb{C})$.

$$\begin{aligned} B_4^{\gamma}(j_1, j_2, j_3, j_4, l_1, l_2, l_3, l_4; i, k) \\ = \sum_{p_f} \begin{pmatrix} l_1 & l_2 & l_3 & l_4 \\ p_1 & p_2 & p_3 & p_4 \end{pmatrix}^{(k)} \left(\int_0^{\infty} dr \frac{1}{4\pi} \sinh^2 r \otimes_{f=1}^4 d_{l_f j_f p_f}^{\gamma j_f j_f}(r) \right) \begin{pmatrix} j_1 & j_2 & j_3 & j_4 \\ p_1 & p_2 & p_3 & p_4 \end{pmatrix}^{(i)}. \end{aligned} \quad (\text{A5})$$

The expression for $d_{jlm}^{\gamma j j}(r)$ has been written in [21,40]:

$$\begin{aligned} d_{jlp}^{\gamma(j,j)}(r) &= (-1)^{\frac{l-l}{2}} \frac{\Gamma(j+i\gamma j+1)}{|\Gamma(j+i\gamma j+1)|} \frac{\Gamma(l-i\gamma j+1)}{|\Gamma(l-i\gamma j+1)|} \frac{\sqrt{2j+1}\sqrt{2l+1}}{(j+l+1)!} \\ &\times \left[(2j)!(l+j)!(l-j)! \frac{(l+p)!(l-p)!}{(j+p)!(j-p)!} \right]^{1/2} e^{-(j-i\gamma j+p+1)r} \\ &\times \sum_s \frac{(-1)^s e^{-2sr}}{s!(l-j-s)!} {}_2F_1[l+1-i\gamma j, j+p+1+s, j+l+2, 1-e^{-2r}], \end{aligned} \quad (\text{A6})$$

where ${}_2F_1$ is the Gauss hypergeometric function. In (A5) we used a short notation for the $(4jm)$ Wigner symbols. They are the contraction of two Wigner $(3jm)$ symbols (the unique intertwiner of three $SU(2)$ representation) labeled by the virtual spin k :

$$\begin{pmatrix} j_1 & j_2 & j_3 & j_4 \\ p_1 & p_2 & p_3 & p_4 \end{pmatrix}^{(k)} = \sum_{p_k} (-1)^{k-p_k} \begin{pmatrix} j_1 & j_2 & k \\ p_1 & p_2 & p_k \end{pmatrix} \begin{pmatrix} k & j_3 & j_4 \\ -p_k & p_3 & p_4 \end{pmatrix}. \quad (\text{A7})$$

The large quantum number limit of the booster functions (A5) has been studied in [31]. They possess an appealing geometrical interpretation in terms of boosted tetrahedra.

APPENDIX B: MONTE CARLO SUMMATION

Suppose we want to compute the sum

$$S = \sum_{j \in \mathcal{J}} a_j, \quad (\text{B1})$$

where \mathcal{J} is a finite subset of $(\mathbb{N}/2)^F$ and j is a multi-index. So far, we are just muddling the waters with a complicated notation to give a nod to the spin foam application. If $V_{\mathcal{J}} = |\mathcal{J}|$ is the cardinality of the set, we can always map it in the interval of natural numbers between 1 and $V_{\mathcal{J}}$. The sum (B1) is a fancy way to represent a sum over an integer index from 1 to $V_{\mathcal{J}}$.

If $V_{\mathcal{J}}$ is very large, the numerical computation of S can result in a highly resource-hungry task. Therefore, we want to use (discrete) Monte Carlo techniques to approximate its value.

We define a uniform probability density function over the set \mathcal{J} as $R_{\mathcal{J}}$. We assume that $R_{\mathcal{J}}$ is normalized to 1. The probability associated to every element of \mathcal{J} with $R_{\mathcal{J}}$ is $1/V_{\mathcal{J}}$. The hypothesis that the set \mathcal{J} is finite ensures that $R_{\mathcal{J}}$ exists.

The fundamental step toward the implementation of Monte Carlo is to interpret the sum (B1) as the expectation value of the terms of the sum a_j using $R_{\mathcal{J}}$

$$S = V_{\mathcal{J}} \sum_{j \in \mathcal{J}} \frac{a_j}{V_{\mathcal{J}}} = V_{\mathcal{J}} E[a(R_{\mathcal{J}})]. \quad (\text{B2})$$

We use a discrete uniform probability distribution because we assume we do not know in advance which term of the sum a_j is contributing the most. This information would allow sampling from a more efficient probability distribution using importance sampling Monte Carlo. Alternatively, it would be possible to use Markov Chain Monte Carlo methods, such as the Metropolis-Hastings algorithm. This approach has been applied recently in [27] to compute observables in spin foams with many boundary degrees of freedom.

We approximate the expectation value in (B2) using a sample of the set \mathcal{J} . We use the probability distribution $R_{\mathcal{J}}$ and randomly extract N_{mc} elements from the set \mathcal{J} . We

denote this set as \mathcal{J}^{mc} . Strictly speaking, \mathcal{J}^{mc} is not a subset of \mathcal{J} since it can contain elements more than once. The Monte Carlo estimate of S is given by

$$S^{mc} = \frac{V_{\mathcal{J}}}{N_{mc}} \sum_{j \in \mathcal{J}^{mc}} a_j. \quad (\text{B3})$$

The law of large numbers ensures that the average of a large number of samples becomes closer and closer to the expected value as more samples are performed. Since S^{mc} is a sample of the sum S we have that

$$S^{mc} \xrightarrow{N_{mc} \rightarrow \infty} S. \quad (\text{B4})$$

The sum S^{mc} (B3) is the Monte Carlo estimate of S . We consider a ‘‘large’’ number of samples N_{mc} as we are interested in a numerical approximation of (B2).

The amount of computational resources necessary to compute S^{mc} scales with the number of samples N_{mc} and not the size of the original set $V_{\mathcal{J}}$. The number of samples N_{mc} is a parameter of the calculation that we can tune. To have a good approximation and save resources, we must find a balanced value for N_{mc} .

APPENDIX C: AITKEN EXTRAPOLATION

The Aitken’s delta-squared process or Aitken extrapolation [41] is a numerical recipe used to accelerate the rate of convergence of a sequence. Suppose you have a convergent sequence S_n with $S = \lim_{n \rightarrow \infty} S_n$. If S_n converges linearly, namely

$$\lim_{n \rightarrow \infty} \frac{|S_n - S|}{|S_{n-1} - S|} = \lambda, \quad (\text{C1})$$

with $0 < \lambda < 1$. Linear convergence means the sequence is closer to its limit by almost the same amount with every step. In this case, we can approximate S starting from the approximate relation

$$\frac{S_n - S}{S_{n-1} - S} \approx \frac{S_{n-1} - S}{S_{n-2} - S}, \quad (\text{C2})$$

and solve for S to find

$$S \approx A[S_n] \equiv \frac{S_n S_{n-2} - S_{n-1}^2}{S_n - 2S_{n-1} + S_{n-2}}. \quad (\text{C3})$$

The sequence $A[S_n]$ is the Aitken extrapolation of S_n and converge to S faster than linear (see the original paper [41] or the book [42] for a proof) meaning that

$$\lim_{n \rightarrow \infty} \frac{|A[S_n] - S|}{|S_n - S|} = 0. \quad (\text{C4})$$

The requirement of linear convergence of S_n is equivalent to asking that for n large enough

$$S_n \approx S + C\lambda^n, \quad (\text{C5})$$

for some constant C and for $|\lambda| < 1$. This allows us to estimate λ without knowing the limit S , as is often the case. In practice, if we knew the value of the limit S , we would not need to extrapolate S_n . The limit of differences

$$\lim_{n \rightarrow \infty} \frac{S_n - S_{n-1}}{S_{n-1} - S_{n-2}} = \lambda. \quad (\text{C6})$$

If the rate convergence of S_n is of higher order (quadratic or more), we do not need to extrapolate as the convergence is already very fast. However, if the convergence is sublinear, we should look for a different extrapolation technique. We refer to the Appendix of [23] for an explicit and simple example.

APPENDIX D: DIAGRAMS OF THE MELONIC SELF-ENERGY AND VERTEX RENORMALIZATION SPIN FOAM AMPLITUDES

In this appendix, we report the wiring diagrams of the self-energy and vertex renormalization spin foam amplitudes, highlighting the combinatorics of the internal faces. These are shown in Figs. 20 and 21. The notation for each intertwiner is used in Eqs. (10) and (11). In order not to clutter the picture, we do not explicitly label the spins.




FIG. 20. Wiring of the 2-complex of the self-energy spin foam diagram. The internal faces are highlighted with different colors. Boundary intertwiners have a blue box and also correspond to the integrals removed to regularize the amplitude.




FIG. 21. Wiring of the 2-complex of the vertex renormalization spin foam diagram. The internal faces are highlighted with different colors. Boundary intertwiners have a blue box and also correspond to the integrals removed to regularize the amplitude.

- [1] C. Rovelli and F. Vidotto, *Covariant Loop Quantum Gravity: An Elementary Introduction to Quantum Gravity and Spinfoam Theory*, Cambridge Monographs on Mathematical Physics (Cambridge University Press, Cambridge, England, 2014).
- [2] A. Perez, The spin foam approach to quantum gravity, *Living Rev. Relativity* **16**, 3 (2013).
- [3] J. Engle, E. Livine, R. Pereira, and C. Rovelli, LQG vertex with finite Immirzi parameter, *Nucl. Phys.* **B799**, 136 (2008).
- [4] L. Freidel and K. Krasnov, A new spin foam model for 4D gravity, *Classical Quantum Gravity* **25**, 125018 (2008).
- [5] W. Kaminski, M. Kisielowski, and J. Lewandowski, Spinfoams for all loop quantum gravity, *Classical Quantum Gravity* **27**, 095006 (2010); **29**, 049502(E) (2012).
- [6] M. Han, Cosmological constant in LQG vertex amplitude, *Phys. Rev. D* **84**, 064010 (2011).
- [7] M. Han, Four-dimensional spinfoam quantum gravity with a cosmological constant: Finiteness and semiclassical limit, *Phys. Rev. D* **104**, 104035 (2021).
- [8] F. Conrady and J. Hnybida, A spin foam model for general Lorentzian 4-geometries, *Classical Quantum Gravity* **27**, 185011 (2010).
- [9] J. W. Barrett, R. J. Dowdall, W. J. Fairbairn, F. Hellmann, and R. Pereira, Lorentzian spin foam amplitudes: Graphical calculus and asymptotics, *Classical Quantum Gravity* **27**, 165009 (2010).
- [10] P. Dona and S. Speziale, Asymptotics of lowest unitary $SL(2, \mathbb{C})$ invariants on graphs, *Phys. Rev. D* **102**, 086016 (2020).
- [11] P. Donà, F. Gozzini, and G. Sarno, Numerical analysis of spin foam dynamics and the flatness problem, *Phys. Rev. D* **102**, 106003 (2020).
- [12] M. Han, Z. Huang, H. Liu, and D. Qu, Complex critical points and curved geometries in four-dimensional

- Lorentzian spinfoam quantum gravity, *Phys. Rev. D* **106**, 044005 (2022).
- [13] J. Engle and C. Rovelli, The accidental flatness constraint does not mean a wrong classical limit, *Classical Quantum Gravity* **39**, 117001 (2022).
- [14] P. Dona, Geometry from local flatness in Lorentzian spin foam theories, *Phys. Rev. D* **107**, 066011 (2023).
- [15] M. Han, Z. Huang, H. Liu, D. Qu, and Y. Wan, Spinfoam on a Lefschetz thimble: Markov chain Monte Carlo computation of a Lorentzian spinfoam propagator, *Phys. Rev. D* **103**, 084026 (2021).
- [16] M. Han, H. Liu, and D. Qu, Complex critical points in Lorentzian spinfoam quantum gravity: 4-simplex amplitude and effective dynamics on double- Δ_3 complex (2023).
- [17] S. K. Asante, B. Dittrich, and H. M. Haggard, Discrete gravity dynamics from effective spin foams, *Classical Quantum Gravity* **38**, 145023 (2021).
- [18] S. K. Asante, B. Dittrich, and J. Padua-Arguelles, Effective spin foam models for Lorentzian quantum gravity, *Classical Quantum Gravity* **38**, 195002 (2021).
- [19] P. Dona and G. Sarno, Numerical methods for EPRL spin foam transition amplitudes and Lorentzian recoupling theory, *Gen. Relativ. Gravit.* **50**, 127 (2018).
- [20] F. Gozzini, A high-performance code for EPRL spin foam amplitudes, *Classical Quantum Gravity* **38**, 225010 (2021).
- [21] S. Speziale, Boosting Wigner's nj-symbols, *J. Math. Phys. (N.Y.)* **58**, 032501 (2017).
- [22] P. Donà, M. Fanizza, G. Sarno, and S. Speziale, Numerical study of the Lorentzian Engle-Pereira-Rovelli-Livine spin foam amplitude, *Phys. Rev. D* **100**, 106003 (2019).
- [23] P. Dona and P. Frisoni, How-to compute EPRL spin foam amplitudes, *Universe* **8**, 208 (2022).
- [24] P. Frisoni, F. Gozzini, and F. Vidotto, Numerical analysis of the self-energy in covariant loop quantum gravity, *Phys. Rev. D* **105**, 106018 (2022).
- [25] P. Donà, P. Frisoni, and E. Wilson-Ewing, Radiative corrections to the Lorentzian Engle-Pereira-Rovelli-Livine spin foam propagator, *Phys. Rev. D* **106**, 066022 (2022).
- [26] P. Frisoni, A numerical approach to the Black-to-White hole transition, [arXiv:2304.02691](https://arxiv.org/abs/2304.02691).
- [27] P. Frisoni, F. Vidotto, and F. Gozzini, Markov chain Monte Carlo methods for graph refinement in spinfoam cosmology, *Classical Quantum Gravity* **40**, 105001 (2023).
- [28] S. K. Asante, J. D. Simão, and S. Steinhaus, Spin-foams as semi-classical vertices: Gluing constraints and a hybrid algorithm, *Phys. Rev. D* **107**, 046002 (2023).
- [29] P. Frisoni, Studying the EPRL spinfoam self-energy, in *16th Marcel Grossmann Meeting on Recent Developments in Theoretical and Experimental General Relativity, Astrophysics and Relativistic Field Theories* (World Scientific, Singapore, 2021), [10.1142/9789811269776_0336](https://doi.org/10.1142/9789811269776_0336).
- [30] P. Frisoni and P. Dona, Summing bulk quantum numbers with Monte Carlo in spin foam theories, [GitHub repository](https://github.com/PietropaoloFrisoni/Monte_Carlo_spinfoams) (2022), https://github.com/PietropaoloFrisoni/Monte_Carlo_spinfoams.
- [31] P. Dona, M. Fanizza, P. Martin-Dussaud, and S. Speziale, Asymptotics of $SL(2, \mathbb{C})$ coherent invariant tensors, *Commun. Math. Phys.* **389**, 399 (2022).
- [32] E. Bianchi, D. Regoli, and C. Rovelli, Face amplitude of spinfoam quantum gravity, *Classical Quantum Gravity* **27**, 185009 (2010).
- [33] P. Dona, M. Han, and H. Liu, Spin-foams and high performance computing, in *Handbook of Quantum Gravity* (Springer, 2023).
- [34] P. Donà, Infrared divergences in the EPRL-FK spin foam model, *Classical Quantum Gravity* **35**, 175019 (2018).
- [35] L. Freidel and D. Louapre, Diffeomorphisms and spin foam models, *Nucl. Phys.* **B662**, 279 (2003).
- [36] A. Riello, Self-energy of the Lorentzian Engle-Pereira-Rovelli-Livine and Freidel-Krasnov model of quantum gravity, *Phys. Rev. D* **88**, 024011 (2013).
- [37] C. Perini, C. Rovelli, and S. Speziale, Self-energy and vertex radiative corrections in LQG, *Phys. Lett. B* **682**, 78 (2009).
- [38] www.computeCanada.ca.
- [39] A. P. Yutsis, I. B. Levinson, and V. V. Vanagas, *Mathematical Apparatus of the Theory of Angular Momentum* (Israel Program for Scientific Translation, Jerusalem, Israel, 1962).
- [40] W. Ruhl, *The Lorentz Group and Harmonic Analysis* (W. A. Benjamin, Inc., New York, 1970).
- [41] A. C. Aitken, XXV.—On Bernoulli's numerical solution of algebraic equations, *Proc. R. Soc. Edinburgh* **46**, 289 (1927).
- [42] A. Sidi, *Practical Extrapolation Methods: Theory and Applications*, Cambridge Monographs on Applied and Computational Mathematics (Cambridge University Press, Cambridge, England, 2003).