3 On-shell recursion relations at tree-level

Recursion relations provide a method for building higher-point amplitudes from lower-point information. The 1988 *off-shell* recursion relations by Berends and Giele [12] construct n-point parton amplitudes from building blocks that are lower-point amplitudes with one leg off-shell (see the review articles [1, 3]). This off-shell method continues to be useful as an algorithm for efficient numerical evaluation of scattering amplitudes. In this review, we focus on the "modern" recursive methods whose building blocks are *on-shell* amplitudes. These *on-shell recursion relations* are elegant in that they use input only from gauge-invariant objects and they have proven very powerful for elucidating the mathematical structure of scattering amplitudes.

In the on-shell approaches, a key idea is to use the power of complex analysis to exploit the analytic properties of on-shell scattering amplitudes. The derivation of on-shell recursion relations is a great example of this, as we shall see in Section 3.1. The most famous on-shell recursion relations are the ***BCFW recursion relations*** of Britto, Cachazo, Feng, and Witten [13, 14], but there are other versions based on the same idea as BCFW, namely the use of complex deformations of the external momenta. We describe the idea here, first in a very general formulation in Section 3.1, then specialize the results to derive the BCFW recursion relations in Section 3.2. We illustrate the BCFW method with a selection of examples, including an inductive proof of the Parke–Taylor formula (2.83). Section 3.3 contains a discussion of when to expect existence of recursion relations in general local QFTs. Finally, in Section 3.4, we present the CSW construction (Cachazo-Svrcek-Witten [15]), also called the ***MHV vertex expansion***.

3.1 Complex shifts and Cauchy's theorem

An on-shell amplitude A_n is characterized by the momenta of the external particles and their type (for example a helicity label h_i for massless particles). We focus here on massless particles so $p_i^2 = 0$ for all $i = 1, 2, \ldots, n$. Of course, momentum conservation $\sum_{i=1}^n p_i^\mu = 0$ is also imposed.

Let us now introduce n complex-valued vectors r_i^μ (some of which may be zero) such that

(i) $\displaystyle\sum_{i=1}^{n} r_i^\mu = 0\,,$

(ii) $r_i \cdot r_j = 0$ for all $i, j = 1, 2, \ldots, n$. In particular $r_i^2 = 0$, and

(iii) $p_i \cdot r_i = 0$ for each i (no sum).

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 01:07:12, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/CBO9781107706620.005

These vectors r_i are used to define n shifted momenta

$$\hat{p}_i^\mu \equiv p_i^\mu + z\, r_i^\mu \qquad \text{with} \quad z \in \mathbb{C}\,. \tag{3.1}$$

Note that

(A) By property (i), momentum conservation holds for the shifted momenta: $\sum_{i=1}^n \hat{p}_i^\mu = 0$.
(B) By (ii) and (iii), we have $\hat{p}_i^2 = 0$, so each shifted momentum is on-shell.
(C) For a non-trivial[1] subset of generic momenta $\{p_i\}_{i\in I}$, define $P_I^\mu = \sum_{i\in I} p_i^\mu$. Then $\hat{P}_I^2$ is *linear* in z:

$$\hat{P}_I^2 = \left(\sum_{i\in I} \hat{p}_i\right)^2 = P_I^2 + z\, 2\, P_I \cdot R_I \qquad \text{with} \quad R_I = \sum_{i\in I} r_i\,, \tag{3.2}$$

because the z^2 term vanishes by property (ii). We can write

$$\hat{P}_I^2 = -\frac{P_I^2}{z_I}(z - z_I) \qquad \text{with} \quad z_I = -\frac{P_I^2}{2P_I \cdot R_I}\,. \tag{3.3}$$

As a result of (A) and (B), we can consider our amplitude A_n in terms of the shifted momenta $\hat{p}_i^\mu$ instead of the original momenta p_i^μ. In particular, it is useful to study the *shifted amplitude* as a function of z; by construction, it is a holomorphic function $\hat{A}_n(z)$. The amplitude with unshifted momenta p_i^μ is obtained by setting $z = 0$, $A_n = \hat{A}_n(z = 0)$.

We specialize to the case where A_n is a **tree-level amplitude**. In that case, the analytic structure of $\hat{A}_n(z)$ is very simple. As discussed in Section 2.8, the tree amplitude does not have any branch cuts – there are no logs, square-roots, etc. at tree-level. The tree amplitude is a rational function of the kinematic variables. As such, its analytic structure is captured by its poles and these are determined by the exchanges of physical particles. Thus, the shifted amplitude $\hat{A}_n(z)$ is a rational function of z and, for generic external momenta, its poles are all simple poles in the z-plane. To see that the poles are simple, consider the Feynman diagrams: the only way we can get poles is from the shifted propagators $1/\hat{P}_I^2$, where $\hat{P}_I$ is a sum of a non-trivial subset of the shifted momenta. By (C) above, $1/\hat{P}_I^2$ gives a simple pole at z_I, and for generic momenta, $P_I^2 \neq 0$ so $z_I \neq 0$. For generic momenta, no Feynman tree diagram can have more than one power of a given propagator $1/\hat{P}_I^2$; and poles of different propagators are at different locations in the z-plane. Hence, for generic momenta, $\hat{A}_n(z)$ only has simple poles and they are all located away from the origin $z = 0$. Note the underlying assumption of locality, i.e. that the amplitudes can be derived from some local Lagrangian so that the propagators determine the poles.

Let us now study the holomorphic function $\frac{\hat{A}_n(z)}{z}$ in the complex z-plane. Pick a contour that surrounds the simple pole at the origin. The residue at this pole is nothing but the unshifted amplitude, $A_n = \hat{A}_n(z = 0)$. Deforming the contour to surround all the other

[1] Non-trivial means at least two and no more than $n-2$ momenta such that $P_I^2 \neq 0$.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 01:07:12, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/CBO9781107706620.005

poles, Cauchy's theorem tells us that

$$A_n = -\sum_{z_I} \mathrm{Res}_{z=z_I} \frac{\hat{A}_n(z)}{z} + B_n \,, \tag{3.4}$$

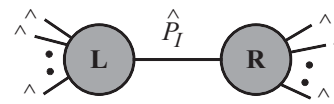
where B_n is the residue of the pole at $z = \infty$. By taking $z \to 1/w$ it is easily seen that B_n is the $O(z^0)$ term in the $z \to \infty$ expansion of A_n.

Now, so what? Well, at a z_I-pole the propagator $1/\hat{P}_I^2$ goes on-shell. In that limit, the shifted amplitude *factorizes* into two on-shell parts,

$$\hat{A}_n(z) \xrightarrow{z \text{ near } z_I} \hat{A}_{\mathrm{L}}(z_I) \frac{1}{\hat{P}_I^2} \hat{A}_{\mathrm{R}}(z_I) \;=\; -\frac{z_I}{z - z_I} \hat{A}_{\mathrm{L}}(z_I) \frac{1}{P_I^2} \hat{A}_{\mathrm{R}}(z_I) \,. \tag{3.5}$$

In the second step we used (3.3). This makes it easy to evaluate the residue at $z = z_I$:

$$-\mathrm{Res}_{z=z_I} \frac{\hat{A}_n(z)}{z} \;=\; \hat{A}_{\mathrm{L}}(z_I) \frac{1}{P_I^2} \hat{A}_{\mathrm{R}}(z_I) \;=$$



. (3.6)

Note that – as opposed to Feynman diagrams – the momentum of the internal line in (3.6) is on-shell, $\hat{P}_I^2 = 0$, and the vertex-blobs represent shifted *on-shell amplitudes* evaluated at $z = z_I$; we call them *subamplitudes*. The rule for the internal line in the diagrammatic representation (3.6) is to write the scalar propagator $1/P_I^2$ of the *unshifted* momenta. Each subamplitude necessarily involves fewer than n external particles, hence all the residues at finite z can be determined in terms of on-shell amplitudes with less than n particles. This is the basis of the recursion relations.

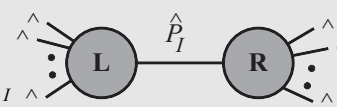
The contribution B_n from the pole at infinity has no similar general expression in terms of lower-point amplitudes; there has been some work on how to compute B_n systematically (see for example [16, 17]), but currently there is not a general constructive method. Thus, in most applications, one assumes – or, much preferably, proves – that $B_n = 0$. This is most often justified by demonstrating the stronger statement that

$$\hat{A}_n(z) \to 0 \quad \text{for} \quad z \to \infty \,. \tag{3.7}$$

If (3.7) holds, we say that the shift (3.1) is ***valid*** (or ***good***).

For a valid shift, the n-point on-shell amplitude is completely determined in terms of lower-point on-shell amplitudes as

$$A_n = \sum_{\text{diagrams } I} \hat{A}_{\mathrm{L}}(z_I) \frac{1}{P_I^2} \hat{A}_{\mathrm{R}}(z_I) = \sum_{\text{diagrams } I}$$



. (3.8)

The sum is over all possible factorization channels I. There is also implicitly a sum over all possible on-shell particle states that can be exchanged on the internal line: for example, for a gluon we have to sum the possible helicity assignments.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 01:07:12, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/CBO9781107706620.005

The recursive formula (3.8) gives a manifestly gauge-invariant construction of scattering amplitudes. Thus (3.8) is the general "prototype" of the ***on-shell recursion relations for tree-level amplitudes*** under a valid shift of the external momenta. We did not use any special properties of 4d spacetime, so the general derivation of the recursion relations is valid in D spacetime dimensions. In the following, we specialize to $D = 4$ again.

3.2 BCFW recursion relations

We shifted all external momenta democratically in (3.1), but with a parenthetical remark that some of the light-like shift-vectors r_i^μ might be trivial, $r_i^\mu = 0$. The BCFW shift is one in which exactly two lines, say i and j, are selected as the only ones with non-vanishing shift-vectors. In $D = 4$ spacetime dimension, the shift is implemented on angle and square spinors of the two chosen momenta:

$$|\hat{i}] = |i] + z\,|j]\,, \qquad |\hat{j}] = |j]\,, \qquad |\hat{i}\rangle = |i\rangle\,, \qquad |\hat{j}\rangle = |j\rangle - z|i\rangle\,. \tag{3.9}$$

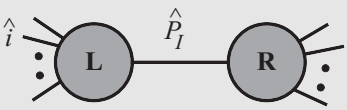
No other spinors are shifted. We call this a $[i, j\rangle$-shift. Note that $[\hat{i}k]$ and $\langle\hat{j}k\rangle$ are linear in z for $k \neq i, j$ while $\langle\hat{i}\,\hat{j}\rangle = \langle ij\rangle$, $[\hat{i}\,\hat{j}] = [ij]$, $\langle\hat{i}k\rangle = \langle ik\rangle$, and $[\hat{j}k] = [jk]$ remain unshifted.

▶ **Exercise 3.1**

Use (2.15) to calculate the shift vectors r_i^μ and r_j^μ corresponding to the shift (3.9). Then show that your shift vectors satisfy properties (i)–(iii) of Section 3.1.

With the two momenta i and j shifted according to (3.9), the BCFW recursion relation for tree amplitudes takes the form

$$A_n = \sum_{\text{diagrams } I} \hat{A}_{\rm L}(z_I)\,\frac{1}{P_I^2}\,\hat{A}_{\rm R}(z_I) = \sum_{\text{diagrams } I}$$



(3.10)

The sum is over all channels I such that the shifted lines i and j are on opposite sides of the factorization diagram in (3.10). As in the general recursion relations (3.8), there is also an implicit sum over all possible on-shell particle states that can be exchanged on the internal line.

Before diving into applications of the BCFW recursion relations (such as proving the Parke–Taylor formula), let us study the shifts a little further. As an example, consider the Parke–Taylor amplitude

$$A_n\big[1^-2^-3^+\ldots n^+\big] \;=\; \frac{\langle 12\rangle^4}{\langle 12\rangle\langle 23\rangle\cdots\langle n1\rangle}\,. \tag{3.11}$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 01:07:12, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/CBO9781107706620.005

Explore the following properties:

▶ **Exercise 3.2**

Convince yourself that for large-z the amplitude (3.11) falls off as $1/z$ under a $[-,-\rangle$-shift (i.e. choose i and j to be the two negative helicity lines). What happens under the three other types of shifts? Note the difference between shifting adjacent/non-adjacent lines.[2]

▶ **Exercise 3.3**

Consider the action of a $[1,2\rangle$-shift of (3.11). Identify the simple pole. Calculate the residue of $\hat{A}_n(z)/z$ at this pole. Compare with (3.4). What happens if you try to repeat this for a $[1,3\rangle$-shift?

▶ **Exercise 3.4**

The spin-1 polarization vectors (2.50) have denominators with $\langle qp\rangle$ and $[qp]$ (with q the reference spinor) which may shift under a BCFW shift (3.9) involving p. Why are there no terms in the on-shell recursion relations (3.8) corresponding to poles at $\langle q\,\hat{p}\rangle = 0$ or $[q\,\hat{p}] = 0$?

The validity of the BCFW recursion relations requires that the boundary term B_n in (3.4) is absent. The typical approach is to show that the shifted amplitude vanishes in the limit of large z, as in (3.7):

$$\hat{A}_n(z) \to 0 \quad \text{for} \quad z \to \infty. \tag{3.12}$$

In pure Yang–Mills theory, an argument [18] based on the background field method establishes the following large-z behavior of color-ordered gluon tree amplitudes under a BCFW shift of adjacent gluon lines i and j of helicity as indicated:

$$\begin{array}{cccccc} [i,j\rangle & [-,-\rangle & [-,+\rangle & [+,+\rangle & [+,-\rangle & \\ \hat{A}_n(z) \sim & \frac{1}{z} & \frac{1}{z} & \frac{1}{z} & z^3 & . \end{array} \tag{3.13}$$

If i and j are non-adjacent, one gains an extra power $1/z$ in each case. Thus any one of the three types of shifts $[-,-\rangle$, $[-,+\rangle$, $[+,+\rangle$ gives valid recursion relations for gluon tree amplitudes.

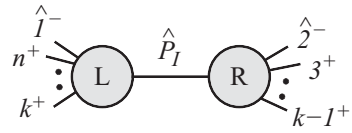
We are now going to use the BCFW recursion relations (3.10) to construct an ***inductive proof of the Parke–Taylor formula*** (3.11). The formula (3.11) is certainly true for $n = 3$, as we saw in Section 2.5, and this establishes the base of the induction. For given n, suppose that (3.11) is true for amplitudes with less than n gluons. Then write down the recursion relation for $A_n[1^-2^-3^+\ldots n^+]$ based on the valid $[1,2\rangle$-shift: adapting from

[2] Of course, we cannot use the large-z behavior of formula (3.11) itself to justify the method to prove this formula! A separate argument is needed and will be discussed shortly.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 01:07:12, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/CBO9781107706620.005

(3.10), we have

$$A_n\big[1^-2^-3^+\ldots n^+\big] = \sum_{k=4}^{n}$$

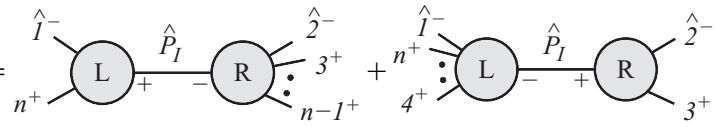


$$= \sum_{k=4}^{n} \sum_{h_I=\pm} \hat{A}_{n-k+3}\big[\hat{1}^-, \hat{P}_I^{h_I}, k^+\ldots, n^+\big]\,\frac{1}{P_I^2}$$
$$\hat{A}_{k-1}\big[-\hat{P}_I^{-h_I}, \hat{2}^-, 3^+\ldots, (k-1)^+\big]\,. \tag{3.14}$$

The internal momentum is labeled P_I, meaning that for a given $k = 4, \ldots, n$ we have $P_I = p_2 + p_3 + \cdots + p_{k-1}$ and $\hat{P}_I = \hat{p}_2 + p_3 + \cdots + p_{k-1}$. There are no diagrams where lines 1 and 2 belong to the same subamplitude, because in that case, the internal momentum would not be shifted and then there is no corresponding residue in (3.4). Only diagrams that preserve the color-ordering of the external states are included. Note that we are also explicitly including the sum over the possible helicity assignments for the particle exchanged on the on-shell internal line: if the exchanged gluon is outgoing from the left subamplitude and has negative helicity, then it will be a positive helicity outgoing gluon as seen from the right subamplitude.

Since one-minus amplitudes $A_n[-+\cdots+]$ vanish except for $n = 3$, (3.14) reduces to

$$A_n\big[1^-2^-3^+\ldots n^+\big] =$$



$$= \hat{A}_3\big[\hat{1}^-, -\hat{P}_{1n}^+, n^+\big]\,\frac{1}{P_{1n}^2}\,\hat{A}_{n-1}\big[\hat{P}_{1n}^-, \hat{2}^-, 3^+\ldots(n-1)^+\big]$$
$$+\,\hat{A}_{n-1}\big[\hat{1}^-, \hat{P}_{23}^-, 4^+\ldots, n^+\big]\,\frac{1}{P_{23}^2}\,\hat{A}_3\big[-\hat{P}_{23}^+, \hat{2}^-, 3^+\big]\,. \tag{3.15}$$

It is understood here that $\hat{P}_I$ is evaluated at the residue value of $z = z_I$ such that $\hat{P}_I^2 = 0$. The notation P_{ij} means $P_{ij} = p_i + p_j$.

The next step is to implement ***special kinematics*** for the 3-point subamplitudes. In the first diagram of (3.15), we have a 3-point anti-MHV amplitude

$$\hat{A}_3\big[\hat{1}^-, -\hat{P}_{1n}^+, n^+\big] \;=\; \frac{[\hat{P}_{1n}\,n]^3}{[n\,\hat{1}][\hat{1}\,\hat{P}_{1n}]}\,. \tag{3.16}$$

Here we use the following convention for analytic continuation:

$$|-p\rangle = -|p\rangle\,, \qquad |-p] = +|p]\,. \tag{3.17}$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 01:07:12, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/CBO9781107706620.005

Since $\hat{P}_{1n}^{\mu} = \hat{p}_1^{\mu} + p_n^{\mu}$, the on-shell condition is

$$0 = \hat{P}_{1n}^2 = 2\hat{p}_1 \cdot p_n = \langle \hat{1} n\rangle[\hat{1} n] = \langle 1n\rangle[\hat{1} n]\,. \tag{3.18}$$

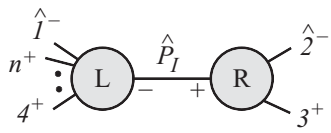
For generic momenta, the only way for the RHS to vanish is for z to take a value such that $[\hat{1} n] = 0$. This means that the denominator in (3.16) vanishes! But so does the numerator: from

$$|\hat{P}_{1n}\rangle[\hat{P}_{1n}\, n] = -\hat{P}_{1n}|n] = -(\hat{p}_1 + p_n)|n] = |1\rangle[\hat{1} n] \;=\; 0\,, \tag{3.19}$$

we conclude that $[\hat{P}_{1n} n] = 0$ since $|\hat{P}_{1n}\rangle$ is not zero. Similarly, one can show that $[\hat{1}\,\hat{P}_{1n}] = 0$. Thus, in the limit of imposing momentum conservation, all spinor products in (3.16) vanish; with the three powers in the numerator versus the two in the denominator, we conclude that special 3-point kinematics force $\hat{A}_3[\hat{1}^-, \hat{P}_{1n}^+, n^+] = 0$. So the contribution from the first diagram in (3.15) vanishes.

In the second diagram of (3.15), the 3-point subamplitude is also anti-MHV, but it does not vanish, since the shift of line 2 is on the angle spinor, not the square spinor. This way, the big abstract recursion formula (3.10) reduces – for the case of the $[-,-\rangle$ BCFW shift of an MHV gluon tree amplitude – to an expression with just a single non-vanishing diagram:

$$A_n\big[1^-2^-3^+\ldots n^+\big] \;=\;$$



$$= \hat{A}_{n-1}\big[\hat{1}^-, \hat{P}_{23}^-, 4^+, \ldots, n^+\big]\,\frac{1}{P_{23}^2}\,\hat{A}_3\big[-\hat{P}_{23}^+, \hat{2}^-, 3^+\big]\,. \tag{3.20}$$

Our inductive assumption is that (3.11) holds for $(n-1)$-point amplitudes. That, together with the result (2.81) for the 3-point anti-MHV amplitude, gives

$$A_n\big[1^-2^-3^+\ldots n^+\big] \;=\; \frac{\langle \hat{1}\,\hat{P}_{23}\rangle^4}{\langle \hat{1}\,\hat{P}_{23}\rangle\langle \hat{P}_{23}\,4\rangle\langle 45\rangle\ldots\langle n\hat{1}\rangle} \times \frac{1}{\langle 23\rangle[23]} \times \frac{[3\,\hat{P}_{23}]^3}{[\hat{P}_{23}\,\hat{2}][\hat{2}3]}\,. \tag{3.21}$$

We could now proceed to evaluate the angle and square spinors for the shifted momenta. But it is more fun to introduce you to a nice little trick. Combine the factors from the numerator:

$$\begin{aligned}\langle \hat{1}\,\hat{P}_{23}\rangle[3\,\hat{P}_{23}] &= -\langle \hat{1}\,\hat{P}_{23}\rangle[\hat{P}_{23}\,3] = \langle \hat{1}|\hat{P}_{23}|3] = \langle \hat{1}|(\hat{p}_2 + p_3)|3] = \langle \hat{1}|\hat{p}_2|3] \\ &= -\langle \hat{1}\hat{2}\rangle[\hat{2}3] = -\langle 12\rangle[23]\,.\end{aligned} \tag{3.22}$$

In the last step we used $\langle \hat{1}\hat{2}\rangle = \langle 12\rangle$ and $|\hat{2}] = |2]$. Playing the same game with the factors in the denominator, we find

$$\langle \hat{P}_{23}\,4\rangle[\hat{P}_{23}\,\hat{2}] = \langle 4|\hat{P}_{23}|\hat{2}] = \langle 4|3|2] = -\langle 43\rangle[32] = -\langle 34\rangle[23]\,. \tag{3.23}$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 01:07:12, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/CBO9781107706620.005

Now use (3.22) and (3.23) in (3.21) to find

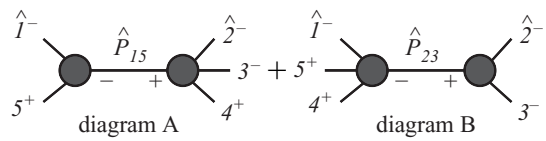
$$A_n\big[1^-2^-3^+\dots n^+\big] = \frac{-\langle 12\rangle^3[23]^3}{\big(-\langle 34\rangle[23]\big)\,\langle 45\rangle\cdots\langle n1\rangle\,\langle 23\rangle[23]\,[23]} = \frac{\langle 12\rangle^4}{\langle 12\rangle\langle 23\rangle\langle 34\rangle\langle 45\rangle\cdots\langle n1\rangle}\,. \tag{3.24}$$

This completes the inductive step. With the 3-point gluon amplitude $A_3\big[1^-2^-3^+\big]$ fixed completely by little group scaling and locality, we have then proven the Parke–Taylor formula for all n. This *is* a lot easier than calculating Feynman diagrams!

You may at this point complain that we have only derived the Parke–Taylor formula recursively for the case where the negative helicity gluons are adjacent. Try your own hands on the proof for the non-adjacent case. In Chapter 4 we will use supersymmetry to derive a more general form of the tree-level gluon amplitudes: it will contain all MHV helicity arrangements in one compact expression.

We have now graduated from MHV-level and are ready to embark on the study of ***NMHV amplitudes***. It is worthwhile to consider the 5-point example $A_5[1^-2^-3^-4^+5^+]$ even though this amplitude is anti-MHV: constructing it with a $[+,+\rangle$-shift is a calculation very similar to the MHV case – and that would by now be boring. So, instead, we are going to use a $[-,-\rangle$-shift to illustrate some of the manipulations used in BCFW recursion.

▷ *Example.* Consider the $[1,2\rangle$-shift recursion relations for $A_5\big[1^-2^-3^-4^+5^+\big]$: there are two diagrams

$A_5\big[1^-2^-3^-4^+5^+\big] =$ 

(3.25)

We have indicated the required helicity for the gluon on the internal line. Had we chosen the opposite helicity option for the internal gluon in diagram A, the righthand subamplitude would have helicity structure $--\,-+$, so it would vanish. Diagram B also vanishes for the opposite choice of the helicity on the internal line. For the helicity choice shown, the righthand subamplitude of diagram B is MHV, $A_3[-\hat{P}^+_{23},\hat{2}^-,3^-]$, and since $|2\rangle$ is shifted, the special 3-particle kinematics actually makes $A_3[-\hat{P}^+_{23},\hat{2}^-,3^-]=0$, just as we saw for the anti-MHV case in the discussion below (3.16). So *diagram B vanishes*, and we can focus on diagram A. Using the Parke–Taylor formula for the MHV subamplitudes, we get

$$A_5\big[1^-2^-3^-4^+5^+\big] = \frac{\langle\hat{1}\hat{P}\rangle^3}{\langle\hat{P}5\rangle\langle 5\hat{1}\rangle}\times\frac{1}{\langle 15\rangle[15]}\times\frac{\langle\hat{2}3\rangle^4}{\langle\hat{2}3\rangle\langle 34\rangle\langle 4\hat{P}\rangle\langle\hat{P}\hat{2}\rangle}\,. \tag{3.26}$$

Here $\hat{P}$ stands for $\hat{P}_{15}=\hat{p}_1+p_5$. We have three powers of $|\hat{P}\rangle$ in the numerator and three in the denominator. A good trick to simplify such expressions is to multiply

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 01:07:12, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/CBO9781107706620.005

(3.26) by $[\hat{P}X]^3/[\hat{P}X]^3$ for some useful choice of X such that $[\hat{P}X] \neq 0$. In this case, it is helpful to pick $X = 2$. Grouping terms conveniently together, we get:

- $\langle \hat{1}\hat{P}\rangle[\hat{P}2] = -\langle \hat{1}|\hat{1}+5|2] = \langle \hat{1}5\rangle[52] = -\langle 15\rangle[25]$ (since $|\hat{1}\rangle = |1\rangle$).
- $\langle 5\hat{P}\rangle[\hat{P}2] = -\langle 5|\hat{1}+5|2] = \langle 51\rangle[\hat{1}2] = \langle 51\rangle[12]$.
- $\langle 4\hat{P}\rangle[\hat{P}2] = -\langle 4|\hat{1}+5|2] = \langle 4|\hat{2}+3+4|2] = \langle 4|3|2] = -\langle 43\rangle[32] = -\langle 34\rangle[23]$.
- $\langle \hat{2}\hat{P}\rangle[\hat{P}2] = -2\,\hat{p}_2\cdot\hat{P} = 2\,\hat{p}_2\cdot(\hat{p}_2+p_3+p_4) = (\hat{p}_2+p_3+p_4)^2-(p_3+p_4)^2$
  $= \hat{P}^2 - \langle 34\rangle[34] = -\langle 34\rangle[34]$,
  since the amplitude is evaluated at z such that $\hat{P}^2 = 0$.

Using these expressions in (3.26) gives

$$A_5\big[1^-2^-3^-4^+5^+\big] \;=\; \frac{[25]^3\langle \hat{2}3\rangle^3}{[12][23][34][15]\langle 34\rangle^3}\,. \tag{3.27}$$

Despite the simplifications, there is some unfinished business for us to deal with: (3.27) depends on the shifted spinors via $\langle \hat{2}3\rangle$. This bracket must be evaluated at the residue value of $z = z_{15}$ which is such that $\hat{P}_{15}^2 = 0$:

$$0 = \hat{P}_{15}^2 = \langle 15\rangle[\hat{1}5]\,, \quad \text{i.e.} \quad 0 = [\hat{1}5] = [15] + z_{15}[25]\,, \quad \text{i.e.} \quad z_{15} = -\frac{[15]}{[25]}\,. \tag{3.28}$$

Use this and momentum conservation to write

$$\langle \hat{2}3\rangle = \langle 23\rangle - z_{15}\langle 13\rangle = \frac{\langle 23\rangle[25]+\langle 13\rangle[15]}{[25]} = \frac{\langle 34\rangle[45]}{[25]}. \tag{3.29}$$

Inserting this result into (3.27) we arrive at the expected anti-MHV Parke–Taylor expression

$$A_5\big[1^-2^-3^-4^+5^+\big] \;=\; \frac{[45]^4}{[12][23][34][45][51]}\,. \tag{3.30}$$

As noted initially, the purpose of this example was not to torture you with a difficult way to compute $A_5[1^-2^-3^-4^+5^+]$. The purpose was to illustrate the methods needed for general cases in a simple context. ◁

You may not be overly impressed with the simplicity of the manipulations needed to reduce the raw output of BCFW. Admittedly it requires some work. If you are unsatisfied, go ahead and try the calculations in this section with Feynman diagrams. Good luck.

Now you have seen the basic tricks needed to manipulate the expressions generated by BCFW. So you should get some exercise.

▶ Exercise 3.5

Let us revisit scalar-QED from the end of Section 2.4. Use little group scaling and locality to determine $A_3(\varphi\,\varphi^*\gamma^\pm)$ and compare with your result from Exercise 2.15. Then use a $[4,\,3\rangle$-shift to show that (see Exercise 2.16)

$$A_4\big(\varphi\,\varphi^*\gamma^+\gamma^-\big) = g^2\,\frac{\langle 14\rangle\langle 24\rangle}{\langle 13\rangle\langle 23\rangle}\,. \tag{3.31}$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 01:07:12, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/CBO9781107706620.005

[Hint: this is not a color-ordered amplitude. See also Exercise 2.16.]
What is the large-z falloff of this amplitude under a $[4, 3\rangle$-shift?[3]

▶ Exercise 3.6

Calculate the 4-graviton amplitude $M_4(1^-2^-3^+4^+)$: first recall that little group scaling & locality fix the 3-graviton amplitudes as in Exercise 2.34. Then employ the $[1, 2\rangle$-shift BCFW recursion relations (they are valid [18, 20]).

Check the little group scaling and Bose-symmetry of your answer for $M_4(1^-2^-3^+4^+)$.

[Hint: your result should match one of the amplitudes in Exercise 2.33.]

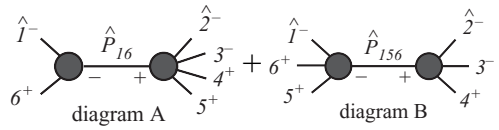
Show that $M_4(1^-2^-3^+4^+)$ obeys the 4-point "KLT relations" [21]

$$M_4(1234) = -s_{12}\, A_4[1234]\, A_4[1243]\,, \tag{3.32}$$

where A_4 is your friend the Parke–Taylor amplitude with negative helicity states 1 and 2, and $s_{12} = -(p_1 + p_2)^2$ is a Mandelstam variable. When you are done, look up ref. [22] to see how difficult it is to do this calculation with Feynman diagrams.

Let us now take a look at some interesting aspects of BCFW for the ***split-helicity*** NMHV amplitude $A_6[1^-2^-3^-4^+5^+6^+]$. Let us first look at the recursion relations following from the $[1, 2\rangle$-shift that we are now quite familiar with. There are two non-vanishing diagrams:

$$A_6\big[1^-2^-3^-4^+5^+6^+\big] \;=\; \text{diagram A} \;+\; \text{diagram B}\,. \tag{3.33}$$



▶ Exercise 3.7

Show that the 23-channel diagram does not contribute in (3.33).

The *first thing* we want to discuss about the 6-gluon amplitude is the 3-particle poles in the expression (3.33). Diagram B involves a propagator $1/P^2_{156}$, so there is a 3-particle pole at $P^2_{156} = 0$. By inspection of the ordering of the external states in $A_6[1^-2^-3^-4^+5^+6^+]$ there should be no distinction between the $(-++)$ 3-particle channels 561 and 345, so we would expect the amplitude to have a pole also at $P^2_{345} = P^2_{126} = 0$. But the $[1, 2\rangle$-shift recursion relation (3.33) does not involve any 126-channel diagram. How can it then possibly encode the correct amplitude? The answer is that it does and that the $P^2_{345} = P^2_{126} = 0$ pole is actually hidden in the denominator factor $\langle\hat{2}\,\hat{P}_{16}\rangle$ of the righthand subamplitude of diagram A in (3.33). Let us show how.

As in the 5-point example (3.26), we multiply the numerator and denominator both by three powers of $[\hat{P}_{16}\,3]$. Then write

$$\langle\hat{2}\,\hat{P}_{16}\rangle[\hat{P}_{16}\,3] = \langle 21\rangle[\hat{1}3] + \langle\hat{2}6\rangle[63]\,. \tag{3.34}$$

[3] Of course, we cannot use the BCFW result for the amplitude to validate the shift, only test self-consistency. The independent demonstration of the large-z falloff under the shift used here is given in [19].

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 01:07:12, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/CBO9781107706620.005

It follows from $\hat{P}_{16}^2 = 0$ that $z_{16} = -[16]/[26]$, and this is then used to show that $\langle \hat{2}6\rangle = (\langle 16\rangle[16] + \langle 26\rangle[26])/[26]$ and $[\hat{1}3] = [12][36]/[26]$. Plug these values into (3.34) to find

$$\langle \hat{2}\hat{P}_{16}\rangle[\hat{P}_{16}\,3] = -\frac{[36]}{[26]}\big(\langle 12\rangle[12] + \langle 16\rangle[16] + \langle 26\rangle[26]\big) = -\frac{[36]}{[26]}\,P_{126}^2\,. \tag{3.35}$$

So there you have it: the 3-particle pole P_{126}^2 is indeed encoded in the BCFW result (3.33).

The *second thing* we want to show you is the actual representation for the 6-gluon NMHV tree amplitude, as it follows from (3.33):

$$A_6\big[1^-2^-3^-4^+5^+6^+\big] = \frac{\langle 3|1+2|6]^3}{P_{126}^2[21][16]\langle 34\rangle\langle 45\rangle\langle 5|1+6|2]} + \frac{\langle 1|5+6|4]^3}{P_{156}^2[23][34]\langle 56\rangle\langle 61\rangle\langle 5|1+6|2]}\,. \tag{3.36}$$

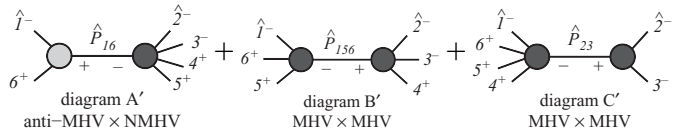
The expression (3.36) may not look quite as delicious as the Parke–Taylor formula, but remember that it contains the same information as the sum of 38 Feynman diagrams!

► **Exercise 3.8**

Check the little group scaling of (3.36). Fill in the details for converting the two diagrams in (3.33) to find (3.36).

The *third thing* we would like to emphasize is that using the $[1,2\rangle$-shift recursion relations is just one way to calculate $A_6[1^-2^-3^-4^+5^+6^+]$. What happens if we use the $[2,1\rangle$-shift? Well, now there are three non-vanishing diagrams:

$$A_6\big[1^-2^-3^-4^+5^+6^+\big] = \quad + \quad + \quad .$$



(3.37)

Special 3-particle kinematics force the diagram A′ to have helicity structure anti-MHV×NMHV, as opposed to the similar diagram A in (3.33) which has to be MHV×MHV. It is the first time we see a lower-point NMHV amplitude show up in the recursion relations. This is quite generic: the BCFW relations are recursive both in particle number n and in N^KMHV level K.

The two BCFW representations (3.33) and (3.37) look quite different. In order for both to describe the same amplitude, there has to be a certain identity that ensures that diagrams A+B = A′ + B′+ C′. To show that this identity holds requires a nauseating trip through Schouten identities and momentum conservation relations in order to manipulate the angle and square brackets into the right form: numerical checks can save you a lot of energy when dealing with amplitudes with more than five external lines. It turns out that the identities that guarantee the equivalence of BCFW expressions such as A+B and A′ + B′+ C′ actually originate from powerful residue theorems [23] related to quite different formulations of the amplitudes. This has to do with the description of amplitudes in the Grassmannian – we get to that in Chapters 9 and 10, but wanted to give you a hint of this interesting point here.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 01:07:12, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/CBO9781107706620.005

▶ **Exercise 3.9**

Show that the BCFW recursion relations based on the $[2,3\rangle$-shift give the following representation of the 6-point "alternating helicity" gluon amplitude:

$$A_6\big[1^+2^-3^+4^-5^+6^-\big] = \{M_2\} + \{M_4\} + \{M_6\}\,, \tag{3.38}$$

where

$$\{M_i\} = \frac{\langle i, i+2\rangle^4 [i+3, i-1]^4}{\tilde{P}_i^2\, \langle i|\tilde{P}_i|i+3]\langle i+2|\tilde{P}_i|i-1]\langle i, i+1\rangle\langle i+1, i+2\rangle[i+3, i-2][i-2, i-1]}\,, \tag{3.39}$$

and $\tilde{P}_i = P_{i,i+1,i+2}$. [Hint: $\{M_4\}$ is the value of the 12-channel diagram.]

In Chapter 9 we discover that each $\{M_i\}$ can be understood as the residue associated with a very interesting contour integral (different from the one used in the BCFW argument).

The *fourth thing* worth discussing further is the poles of scattering amplitudes. Color-ordered tree amplitudes can have physical poles only when the sum of momenta of *adjacent* external lines go on-shell. We touched on this point already in Section 2.8. There we also noted that MHV gluon amplitudes do not have multi-particle poles, only 2-particle poles. Now you have seen that 6-gluon NMHV amplitudes have both 2- and 3-particle poles. But as you stare intensely at (3.36), you will also note that there is a strange denominator-factor $\langle 5|1+6|2]$ in the result from each BCFW diagram. This does not correspond to a physical pole of the scattering amplitude: it is a ***spurious pole***. The residue of this unphysical pole better be zero – and it is: the spurious pole cancels in the sum of the two BCFW diagrams in (3.36). It is typical that BCFW packs the information of the amplitudes into compact expressions, but the cost is the appearance of spurious poles; this means that in the BCFW representation the locality of the underlying field theory is not manifest. Elimination of spurious poles in the representations of amplitudes leads to interesting results [24] that we discuss in Chapter 10.

Finally, for completeness, note that the color-ordered NMHV amplitudes $A_6[1^-2^-3^+4^-5^+6^+]$ and $A_6[1^-2^+3^-4^+5^-6^+]$ are inequivalent to the split-helicity amplitude $A_6[1^-2^-3^-4^+5^+6^+]$. More about this in Chapter 4.

Other comments:

1) In our study of the recursion relations, we kept insisting on "generic" momenta. However, special limits of the external momenta place useful and interesting constraints on the amplitudes: the behavior of amplitudes under ***collinear limits*** and ***soft limits*** is described in Section 2.8.

2) In some cases, the shifted amplitudes have "better than needed" large-z behavior. For example, this is the case for a BCFW shift of two non-adjacent same-helicity lines in the color-ordered Yang–Mills amplitudes: $\hat{A}_n(z) \to 1/z^2$ for large z. The vanishing of the amplitude at large z implies the validity of a recursion relation for A_n. Let us

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 01:07:12, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/CBO9781107706620.005

briefly outline the reason. Start with $\oint_{\mathcal{C}} \frac{\hat{A}_n(z)}{z} = 0$ with $\mathcal{C}$ a contour that surrounds all the simple poles. The unshifted amplitude A_n, which is the residue of the $z = 0$ pole, is therefore (minus) the sum of all the other residues. We write $A_n = \sum_I d_I$, where d_I is the factorization diagram (3.8) associated with the residue at $z = z_I$. This summarizes the derivation of the recursion relations from Section 3.1. Now, an extra power in the large-z falloff, $\hat{A}_n(z) \sim 1/z^2$, means that there is also a ***bonus relation*** from $\oint_{\mathcal{C}} \hat{A}_n(z) = 0$ (with $\mathcal{C}$ as before). It gives $\sum_I d_I z_I = 0$, because there is no residue at $z = 0$. Bonus relations have practical applications, for example they have been used to demonstrate equivalence of different expressions for graviton amplitudes [25].

3.3 When does it work?

In Section 2.6 we learned that 3-point amplitudes for massless particles are uniquely determined by little group scaling, locality, and dimensional analysis. As we have just seen, with the on-shell BCFW recursion relations, we can construct all higher-point gluon tree amplitudes from the input of just the 3-point gluon amplitudes. That is a lot of information obtained from very little input! It prompts suspicion: when can we expect on-shell recursion to work? We will look at some examples now.

Yang–Mills theory and gluon scattering. The Yang–Mills Lagrangian (2.66) has two types of interaction terms, the cubic $A^2\partial A$ and the quartic A^4. Given the former, the latter is needed for gauge invariance, and the quartic must be included along with the cubic as interaction vertices in the Feynman rules. However, existence of valid BCFW recursion relations indicates that the cubic term $A^2\partial A$ fully captures the information needed for all on-shell gluon amplitudes, at least at tree-level, with no need for A^4. The key distinction is that the 3-vertex (and hence the Feynman rule cubic vertex) is an *off-shell* gauge non-invariant object, while the 3-point *on-shell* amplitude is gauge invariant. Since A^4 is determined from $A^2\partial A$ by the requirement of off-shell gauge invariance of the Lagrangian, it contains no new on-shell information. In a sense, that is why the recursion relations for on-shell gluon amplitudes even have a chance to work with input only from the on-shell 3-point amplitudes.

We can rephrase the information contents of $A^2\partial A$ in a more physical way. The actual input is then this: a 4d local theory with massless spin-1 particles (and no other dynamical states) and a dimensionless coupling constant. With valid recursion relations, this information is enough to fix the entire gluon tree-level scattering matrix!

Scalar-QED. As a second example, consider scalar-QED. The interaction between the photons and the scalar particles created/annihilated by a complex scalar field φ is encoded by the covariant derivatives $D_\mu = \partial_\mu - ieA_\mu$ in

$$\mathcal{L} \supset -|D\varphi|^2 = |\partial\varphi|^2 + ieA^\mu\big[(\partial_\mu\varphi^*)\varphi - \varphi^*\partial_\mu\varphi\big] - e^2 A^\mu A_\mu \varphi^*\varphi\,. \tag{3.40}$$

In terms of Feynman diagrams, the tree amplitude $A_4(\varphi\,\varphi^*\gamma\,\gamma)$ is constructed from the sum of two pole diagrams and the contact term from the quartic interaction (Exercise 2.17).

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 01:07:12, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/CBO9781107706620.005

We have seen in Exercise 3.5 that this 4-point amplitude is constructible via BCFW. So it is clear that only the information in the 3-point vertices is needed, and the role of $A_\mu A^\mu \varphi^* \varphi$ is just to ensure off-shell gauge invariance of the Lagrangian. Thus this case is just like the Yang–Mills example above.

Thus emboldened, let us try to compute the 4-scalar tree amplitude $A_4(\varphi\,\varphi^*\varphi\,\varphi^*)$ using BCFW recursion (3.10). Using a $[1,3\rangle$-shift, there are two diagrams and their sum simplifies to

$$A_4^{\text{BCFW}}(\varphi\,\varphi^*\varphi\,\varphi^*) \;=\; \tilde{e}^2\,\frac{\langle 13\rangle^2\langle 24\rangle^2}{\langle 12\rangle\langle 23\rangle\langle 34\rangle\langle 41\rangle}\,. \tag{3.41}$$

If, however, we calculate this amplitude using Feynman rules from the interaction terms in (3.40), we get

$$A_4^{\text{Feynman}}(\varphi\,\varphi^*\varphi\,\varphi^*) \;=\; \tilde{e}^2\left(1+\frac{\langle 13\rangle^2\langle 24\rangle^2}{\langle 12\rangle\langle 23\rangle\langle 34\rangle\langle 41\rangle}\right). \tag{3.42}$$

Ugh! So BCFW did not compute the amplitude we expected. So what did it compute? Well, let us think about the input that BCFW knows about: 4d local theory with massless spin-1 particles and charged massless spin-0 particles (and no other dynamical states) and a dimensionless coupling constant. Note that included in this input is the possibility of a 4-scalar interaction term $\lambda|\varphi|^4$. So more generally, we should consider the scalar-QED action from (2.64):

$$\begin{aligned}\mathcal{L} &= -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} - |D\varphi|^2 - \frac{1}{4}\lambda|\varphi|^4 \\ &= -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} - |\partial\varphi|^2 + ieA^\mu\big[(\partial_\mu\varphi^*)\varphi - \varphi^*\partial_\mu\varphi\big] - e^2A^\mu A_\mu\varphi^*\varphi - \frac{1}{4}\lambda|\varphi|^4\,.\end{aligned} \tag{3.43}$$

In Exercise 2.19 you were asked to calculate $A_4\big(\varphi\,\varphi^*\varphi\,\varphi^*\big)$ in this model. The answer was given in (2.65): it is

$$A_4\big(\varphi\,\varphi^*\varphi\,\varphi^*\big) \;=\; -\lambda + \tilde{e}^2\left(1+\frac{\langle 13\rangle^2\langle 24\rangle^2}{\langle 12\rangle\langle 23\rangle\langle 34\rangle\langle 41\rangle}\right). \tag{3.44}$$

So it is clear now that we have a family of scalar-QED models, labeled by λ, and that our BCFW calculation produced the very special case of $\lambda = \tilde{e}^2$. How can we understand this? Validity of the recursion relations requires the absence of the boundary term B_n (see Section 3.1). For the general family of scalar-QED models, there is a boundary term under the $[1,3\rangle$-shift, and its value is $-\lambda + \tilde{e}^2$ (as can be seen from (3.44) by direct computation). The special choice $\lambda = \tilde{e}^2$ eliminates the boundary term, and that is then what BCFW without a boundary term computes.

The lesson is that for general λ, there is no way the 3-point interactions can know the contents of $\lambda|\varphi|^4$: it provides independent gauge-invariant information. That information needs to be supplied in order for recursion to work, so in this case one can at best expect recursion to work beyond 4-point amplitudes. The exception is of course if some symmetry, or other principle, determines the information in $\lambda|\varphi|^4$ in terms of the 3-field terms. This is what we find for $\lambda = \tilde{e}^2$. Actually, the expression (3.41) is the correct result for certain 4-scalar amplitudes in $\mathcal{N}=2$ and $\mathcal{N}=4$ super Yang–Mills theory (see Chapter 4), and in

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 01:07:12, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/CBO9781107706620.005

those cases the coupling of the 4-scalar contact term *is* fixed by the Yang–Mills coupling by supersymmetry.

Scalar theory $\lambda\phi^4$. The previous example makes us wary of $\lambda\phi^4$ interaction in the context of recursion relations – and rightly so. Suppose we consider $\lambda\phi^4$ theory with no other interactions. It is clear that one piece of input must be given to start any recursive approach, namely in this case the 4-scalar amplitude $A_4 = \lambda$. In principle, one might expect on-shell recursion to determine all tree-level A_n amplitudes with $n > 4$ from $A_4 = \lambda$ – what else could interfere? After all, this is the only interaction in the Feynman diagrams. However, given that the 6-scalar amplitude is $A_6 = \lambda^2(\frac{1}{s_{123}} + \dots)$, it is clear that any BCFW shift gives $O(z^0)$-behavior for large z and hence there are no BCFW recursion relations without boundary term for A_6 in $\lambda\phi^4$ theory. Inspection of the Feynman diagrams reveals that the $O(z^0)$-contributions are exactly the diagrams in which the two shifted lines belong to the same vertex. The sum of such diagrams equals the boundary term B_n from (3.4). Thus, in this case of $\lambda\phi^4$ theory one can reconstruct B_n recursively.[4] Hence A_4 does suffice to completely determine all tree amplitudes A_n for $n > 4$ in $\lambda\phi^4$ theory; but it is (in more than one sense) a rather trivial example.

$\mathcal{N} = 4$ **super Yang–Mills theory.** This is the favorite theory of many amplitugicians. We will review the theory in more detail in Section 4.4, for now we just comment on a few relevant aspects. The spectrum consists of 16 massless states: gluons $g^{\pm}$ of helicity ± 1, four gluinos λ^a and λ_a of helicity $\pm 1/2$, and six scalars S^{ab}. The indices $a, b = 1, 2, 3, 4$ are labels for the global $SU(4)$ symmetry. The Lagrangian contains standard gluon self-interactions, with standard couplings to the gluinos and the scalars. All fields transform in the adjoint of the $SU(N)$ gauge group, so we consider color-ordered tree amplitudes defined in the same way as the color-ordered gluon amplitudes. The Lagrangian includes a scalar 4-point interaction term of a schematic form $[S, S]^2$. It contains, for example, the interaction $S^{12}S^{23}S^{34}S^{41}$. The result for the corresponding color-ordered amplitude is (suppressing the gauge coupling constant):

$$A_4\left[S^{12}S^{23}S^{34}S^{41}\right] = 1. \tag{3.45}$$

Since this amplitude has no poles, it cannot be obtained via direct factorization. Actually, the amplitude (3.45) and its cousin 4-scalar amplitudes with equivalent $SU(4)$ index structures are the only tree amplitudes of $\mathcal{N} = 4$ SYM that cannot be obtained from the BCFW recursion formula (3.10); that may seem surprising, but it is true [27].

When supersymmetry is incorporated into the BCFW recursion relations, *all* tree amplitudes of $\mathcal{N} = 4$ SYM can be determined by the 3-point gluon vertex alone. The so-called super-BCFW shift mixes the external states in such a way that even the 4-scalar amplitude (3.45) can be constructed recursively. We will introduce the super-BCFW shift in Section 4.5.

[4] See [16]. Or avoid the term at infinity by using an all-line shift [26], to be defined in Section 3.4.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 01:07:12, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/CBO9781107706620.005

Gravity. We have already encountered the 4-point MHV amplitude $M_4(1^-2^-3^+4^+)$: you "discovered" it from little group scaling in Exercise 2.33 and constructed it with BCFW in Exercise 3.6. The validity of the BCFW recursion relations for all tree-level graviton amplitudes [18, 20] means that the entire on-shell tree-level S-matrix for gravity is determined completely by the 3-vertex interaction of three gravitons. In contrast, the expansion of the Einstein–Hilbert action $\frac{1}{2\kappa^2}\int d^4x\,\sqrt{-g}R$ around the flatspace Minkowski metric $g_{\mu\nu} = \eta_{\mu\nu} + \kappa\, h_{\mu\nu}$ contains *infinitely* many interaction terms. It is remarkable that all these terms are totally irrelevant from the point of view of the on-shell tree-level S-matrix; their sole purpose is to ensure diffeomorphism invariance of the off-shell Lagrangian. For on-shell (tree) amplitudes, we do not need them. Much more about gravity amplitudes in Chapter 12.

Summary. We have discussed when to expect to have recursion relations for tree-level amplitudes. The main lesson is that we do not get something for nothing: input must be given and we can only expect to recurse that input with standard BCFW when all other information in the theory is fixed by our input via gauge invariance. If another principle – such as supersymmetry – is needed to fix the interactions, then that principle should be incorporated into the recursion relations for a successful recursive approach. Further discussion of these ideas can be found in [26], mostly in the context of another recursive approach known as CSW, which we will discuss briefly next.

3.4 MHV vertex expansion (CSW)

We introduced recursion relations in Section 3.1 in the context of general shifts (3.1) satisfying the set of conditions (i)–(iii). Then we specialized to the BCFW shifts in Section 3.2. Now we would like to show you another kind of recursive structure.

Consider a shift that is implemented via a "holomorphic" square-spinor shift:

$$|\hat{i}] = |i] + z\, c_i |X] \quad \text{and} \quad |\hat{i}\rangle = |i\rangle\,. \tag{3.46}$$

Here $|X]$ is an arbitrary reference spinor and the coefficients c_i satisfy $\sum_{i=1}^{n} c_i |i\rangle = 0$.

▶ **Exercise 3.10**

Show that the square-spinor shift (3.46) gives shift-vectors r_i that fulfill requirements (i)–(iii) in Section 3.1.

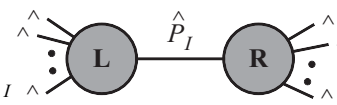
The choice $c_1 = \langle 23\rangle$, $c_2 = \langle 31\rangle$, $c_3 = \langle 12\rangle$, and $c_i = 0$ for $i = 4, \ldots, n$ implies that the shifted momenta satisfy momentum conservation. This particular realization of the square-spinor shift is called the Risager shift [28].

We consider here a situation where all $c_i \neq 0$ so that all momentum lines are shifted via (3.46) – this is an ***all-line shift***. It can be shown [29] that N^KMHV gluon tree amplitudes fall off as $1/z^K$ for large z under all-line shift. So this means that all gluon tree-level

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 01:07:12, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/CBO9781107706620.005

amplitudes can be constructed with the all-line shift recursion relations; except the MHV amplitudes ($K = 0$). It turns out that in this formulation of recursion relations, the tower of MHV amplitudes constitutes the basic building blocks for the N^KMHV amplitudes. Let us see how this works for NMHV. The recursion relations give

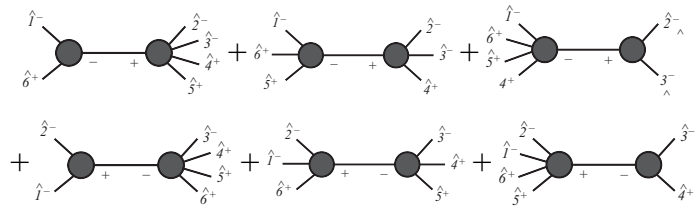
$$A_n^{\mathrm{NMHV}} = \sum_{\text{diagrams } I}$$



. (3.47)

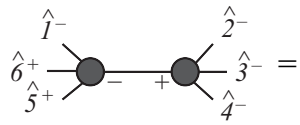
If you consider the possible assignments of helicity labels on the internal line, you will see that there are two options: either the diagram is anti-$\mathrm{MHV}_3 \times$NMHV or MHV$\times$MHV. The former option vanishes by special kinematics of the 3-point anti-MHV vertex, just as in the case of the first diagram in (3.15). So all subamplitudes in (3.47) are MHV. Let us write down the example of the split-helicity NMHV 6-gluon amplitude:

$$A_n\big[1^-2^-3^-4^+5^+6^+\big] =$$



(3.48)

All six diagrams are non-vanishing and this may look a little daunting, especially compared with the BCFW version where there were just two diagrams in the simplest version (3.36). However, the diagrams in (3.48) are easier to evaluate than the BCFW diagrams: the MHV amplitudes depend only on angle spinors, so the only way they know about the square-spinor shift is through the internal line angle spinors $|\hat{P}_I\rangle$, for example



$$= \frac{\langle 1\hat{P}_I\rangle^4}{\langle 1\hat{P}_I\rangle\langle\hat{P}_I 5\rangle\langle 56\rangle\langle 61\rangle}\,\frac{1}{P_{156}^2}\,\frac{\langle 23\rangle^4}{\langle 23\rangle\langle 34\rangle\langle 4\hat{P}_I\rangle\langle\hat{P}_I 2\rangle}\,. \tag{3.49}$$

We can write

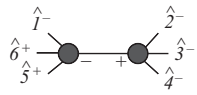
$$|\hat{P}_I\rangle\frac{[\hat{P}_I X]}{[\hat{P}_I X]} = \hat{P}_I|X]\,\frac{1}{[\hat{P}_I X]} = P_I|X]\,\frac{1}{[\hat{P}_I X]}\,. \tag{3.50}$$

In the last step we can drop the hat, because the shift of $(\hat{P}_I)^{\dot{a}b}$ is proportional to the reference spinor $[X|^b$ of the shift (3.46). Note that the diagrams are necessarily invariant under little group scaling associated with the internal line. Therefore the factors $\frac{1}{[\hat{P}_I X]}$ in (3.50) cancel out of each diagram and we can use the prescription

$$|\hat{P}_I\rangle \to P_I|X]\,. \tag{3.51}$$

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 01:07:12, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/CBO9781107706620.005

This gives



$$= \frac{\langle 1|P_{156}|X]^4}{\langle 1|P_{156}|X]\langle 5|P_{156}|X]\langle 56\rangle\langle 61\rangle} \frac{1}{P_{156}^2} \frac{\langle 23\rangle^4}{\langle 23\rangle\langle 34\rangle\langle 4|P_{156}|X]\langle 2|P_{156}|X]} \tag{3.52}$$

and similarly for the other "MHV vertex diagrams" in (3.48). Note that we can drop the indication $\hat{}$ of the shift on the external lines in the MHV vertex diagrams since the square-spinor shift does not affect the MHV vertices and all that is needed is the prescription (3.51) for the internal lines.

In general, each diagram depends explicitly on the reference spinor $|X]$, but of course the full tree amplitude cannot depend on an arbitrary spinor: the Cauchy theorem argument of Section 3.1 guarantees that the sum of all the diagrams will be independent of $|X]$ and reproduce the correct tree amplitude. Numerically, it is not hard to verify independence of $|X]$ and that the expressions (3.48) and (3.36) indeed produce the same scattering amplitude.

The expansion of the amplitude in terms of MHV vertex diagrams generalizes beyond the NMHV level. In general, the N^KMHV tree amplitude is written as a sum of all tree-level diagrams with precisely $K+1$ MHV vertices evaluated via the replacement rule (3.51). This construction of the amplitude is called the ***MHV vertex expansion***: it can be viewed as the closed-form solution to the all-line shift recursion relations. However, it was discovered by Cachazo, Svrcek, and Witten in 2004 [15] before the introduction of recursion relations from complex shifts. The method is therefore also known as the ***CSW expansion*** and the rule (3.51) is called the ***CSW prescription***. The first recursive derivation of the MHV vertex expansion was given by Risager [28] using the 3-line Risager shift mentioned above applied to the three negative helicity line of NMHV amplitudes. The all-line shift formulation was first presented in [29].

▶ **Exercise 3.11**

Construct $A_5[1^-2^-3^-4^+5^+]$ from the CSW expansion. Make a choice for the reference spinor $|X]$ to simplify the calculation and show that the result agrees with the anti-MHV Parke–Taylor formula (2.115).

The MHV vertex expansion was the first construction of gluon amplitudes from on-shell building blocks. The method is valid also in other cases, for example in super Yang–Mills theory [27, 29] or Higgs amplitudes with gluons and partons [30, 31]. There are also applications of the MHV vertex expansion at loop-level – for a review see [32] and references therein.

The MHV vertex expansion can also be derived directly from a Lagrangian [33]: a field redefinition and suitable light-cone gauge choice brings it to a form with an interaction term for each MHV amplitude. The N^KMHV amplitudes are then generated from the MHV vertex Lagrangian by gluing together the MHV vertices. The reference spinor $|X]$ arises from the light-cone gauge choice. There is also a twistor-action formulation of the MHV vertex expansion [34].

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 01:07:12, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/CBO9781107706620.005

In the case of the BCFW shift, we have applied it to gluon as well as graviton amplitudes. A version of the MHV vertex expansion was proposed for gravity in [35] based on the Risager shift. However, the method fails for NMHV amplitudes for $n \geq 12$: under the Risager shift, $\hat{A}_n(z) \sim z^{12-n}$ for large-z, so for $n \geq 12$ there is a boundary term obstructing the recursive formula [36]. An analysis of validity of all-line shift recursion relations in general 4d QFTs can be found in [26].

At this stage, you may wonder why tree-level gluon scattering amplitudes have so many different representations: one from the MHV vertex expansion and other forms arising from BCFW applied to various pairs of external momenta. The CSW and BCFW representations reflect different aspects of the amplitudes, but they turn out to be closely related. We need more tools to learn more about this. So read on.

Downloaded from https://www.cambridge.org/core. UNIVERSIDAD AUTONOMA METROPOLITANA, on 14 Nov 2026 at 01:07:12, subject to the Cambridge Core terms of use, available at https://www.cambridge.org/core/terms. https://doi.org/10.1017/CBO9781107706620.005