1

Introduction

The traditional formulation of quantum field theory—encoded in its very name—is built on the two pillars of *locality* and *unitarity* [2]. The standard apparatus of Lagrangians and path integrals allows us to make these two fundamental principles manifest. This approach, however, requires the introduction of a large amount of unphysical redundancy in our description of physical processes. Even for the simplest case of scalar field theories, there is the freedom to perform field redefinitions. Starting with massless particles of spin 1 or higher, we are forced to introduce even larger, gauge redundancies [2].

Over the past few decades, there has been a growing realization that these redundancies hide amazing physical and mathematical structures lurking within the heart of quantum field theory. This has been seen dramatically at strong coupling in gauge/gauge (see, e.g. [3–5]) and gauge/gravity dualities [6]. The past decade has uncovered further remarkable new structures in field theory even at weak coupling, seen in the properties of scattering amplitudes in gauge theories and gravity (for reviews, see [7–12]). The study of scattering amplitudes is fundamental to our understanding of field theory, and fueled its early development in the hands of Feynman, Dyson, and Schwinger among others. It is therefore surprising to see that even here, by committing so strongly to particular, gauge-redundant descriptions of the physics, the usual formalism is completely blind to astonishingly simple and beautiful properties of the gauge-invariant physical observables of the theory.

Many of the recent developments have been driven by an intensive exploration of $\mathcal{N} = 4$ supersymmetric Yang–Mills (SYM) in the planar limit [12, 13]. The all-loop integrand for scattering amplitudes in this theory can be determined by a generalization of the BCFW recursion relations [14], in a way that is closely tied to remarkable new structures in algebraic geometry, associated with contour integrals over the Grassmannian $G(k,n)$ [15–18]. This makes both the *conformal* and long-hidden *dual conformal* invariance of the theory (which together close into the infinite-dimensional Yangian symmetry) completely manifest [19]. It is remarkable that a single function of external kinematical variables can be interpreted as a scattering amplitude in one space-time, and as a Wilson loop in another (for a review, see [12]). Each of these descriptions makes a commitment to locality in its own space-time, making it impossible to see the dual picture. By contrast, the Grassmannian picture makes no mention of locality or unitarity, and does not commit to any gauge-redundant description of the physics, allowing it to manifest *all* the symmetries of the theory.

There has also been extraordinary progress in determining the amplitude itself beyond the integrand, using the technology of symbols of transcendental functions to powerfully constrain and control the polylogarithms occurring in the final results [20, 21]. While a global picture is still missing, a huge amount of data has been generated. The symbol for all 2-loop MHV amplitudes has been determined [22] (see also [23]), and a handful of 2-loop NMHV and 3-loop MHV symbols have been found [24–26]. Remarkable strategies have also been presented to bootstrap amplitudes to very high loop orders [27–31]. Many of these ideas have a strong resonance with the explosion of progress in the last decade using integrability to find exact results in planar $\mathcal{N}=4$, starting with the spectacular solution of the spectral problem for anomalous dimensions [13, 32].

All of these developments have made it completely clear that there are powerful new mathematical structures underlying the extraordinary properties of scattering amplitudes in gauge theories. If history is any guide, formulating and understanding the physics in a way that makes the symmetries manifest should play a central role in the story. The Grassmannian picture does this, but up to this point there has been little understanding as to why this formulation exists, exactly how it works, or where it comes from physically. Our primary goal in this note is to resolve this unsatisfactory state of affairs.

We will derive the connection between scattering amplitudes and the Grassmannian, starting physically from first principles. This will lead us into direct contact with several beautiful and active areas of current research in mathematics [33–42]. The past few decades have seen vigorous interactions between physics and mathematics in a wide variety of areas, but what is going on here involves *new* areas of mathematics that have only very recently played any role in physics, involving simple but deep ideas ranging from combinatorics to algebraic geometry. It is both startling and exciting that such elementary mathematical notions are found at the heart of the physics of scattering amplitudes.

This new way of thinking about scattering amplitudes involves many novel physical and mathematical ideas. Our presentation will be systematic, and we have endeavored to make it self contained and completely accessible to physicists. While we will discuss a number of mathematical results—some of them new—we will usually be content with the physicist's level of rigor. While the essential ideas here are all very simple, they are tightly interlocking, and range over a wide variety of areas—most of which are unfamiliar to most physicists. Thus, before jumping into the detailed exposition, as a guide to the reader we end this introductory chapter by giving a roadmap of the logical structure and content of the book.

In Chapter 2, we introduce the central physical idea motivating our work, which is to focus on *on-shell diagrams*, obtained by gluing together fundamental 3-particle amplitudes and integrating over the on-shell phase space of internal particles. These objects are of central importance to the understanding scattering amplitudes. We will see that scattering amplitudes in planar $\mathcal{N}=4$—to all loop orders—can be represented *directly* in terms of on-shell processes. In this picture, "virtual particles" make no appearance at all. We should emphasize that we are not merely using on-shell information to determine scattering amplitudes, but rather seeing that the amplitudes can be directly computed in terms of fully on-shell processes. The off-shell, virtual particles familiar from Feynman diagrams are replaced by internal, *on-shell* particles (with generally complex momenta).

In our study of on-shell diagrams, we will see that different diagrams related by certain elementary moves can be physically equivalent, leading to the natural question of how to invariantly characterize their physical content. Remarkably, the invariant content of on-shell diagrams turns out to be characterized by *combinatorial* data. We discuss this in detail in Chapter 3 where we show how a long-known and beautiful connection between permutations and scattering amplitudes in integrable $(1+1)$-dimensional theories generalizes to more realistic theories in $(3+1)$ dimensions.

In Chapter 4 we turn to actually calculating on-shell diagrams and find that the most natural way of carrying out the computations is to associate each diagram with a certain differential form on an *auxiliary* Grassmannian. In Chapter 5 and 6 we show how the invariant, combinatorial content of an on-shell diagram is reflected in the Grassmannian directly. This is described in terms of a surprisingly simple stratification of the configurations of k-dimensional vectors endowed with a cyclic ordering, classified by the linear dependencies among consecutive chains of vectors. For the real Grassmannian, this stratification can be equivalently described in an amazingly simple and beautiful way as nested 'boundaries' of the *positive part* of the Grassmannian [33], which is motivated by the theory of totally positive matrices [34, 43, 44]. Each on-shell diagram can then be associated with a particular configuration or "stratum" among the boundaries of the positive Grassmannian.

In Chapter 7 we make contact with the Grassmannian contour integral of reference [15], which is now seen as a compact way of representing the natural, invariant top-form on the positive Grassmannian. This form of the measure allows us to easily identify the conformal and dual conformal symmetries of the theory, which are related by a simple mapping of permutations described in Chapter 8. In Chapter 9, we show that the invariance of scattering amplitudes under the action of the level-one generators of the Yangian has a transparent interpretation: these generators correspond to the leading nontrivial diffeomorphisms that preserve all the cells of the positive Grassmannian.

In Chapter 10 we begin a systematic classification of Yangian-invariants and their relations by first describing a combinatorial test to determine whether an on-shell diagram has non-vanishing kinematical support (and if it has support, how many solutions exist). In Chapter 11 a geometric basis is given for all the myriad, highly nontrivial identities satisfied among Yangian-invariants. This completes the classification of *all* Yangian-Invariants together with *all* their relations. In Chapter 13, we give a tour of this classification as it emerges through N^4MHV.

In Chapter 14 we show that the story for scattering amplitudes in integrable $(1+1)$-dimensional theories—in particular, the Yang–Baxter relation—can be understood as a special case of our general results regarding on-shell diagrams. We further show that scattering amplitudes for the ABJM theory in $(2+1)$ dimensions [45] can also be computed in terms of a natural specialization of on-shell diagrams: those associated with the null orthogonal Grassmannian. And we initiate the study of on-shell diagrams in theories with less (or no) supersymmetry in Chapter 15.

The positive Grassmannian is naturally endowed with a rich mathematical structure known as a *cluster algebra*—the original theory of which was developed in [35] and has since been generalized to the theory of *cluster varieties* in [37, 38]. Remarkably, this structure has made striking appearances in widely disparate parts of physics in the last decade—from conformal blocks for higher Toda theories [36, 46], to wall-crossing phenomena [47, 48], to quiver gauge theories with $\mathcal{N}=1$ super-conformal symmetry [49–54], to soliton solutions to the KP equation [55–57]. We briefly review this story in Chapter 16, and also summarize its various physical manifestations in hopes of stimulating a deeper understanding for these extremely surprising connections between physics and mathematics.

In Chapter 17 we move beyond the discussion of individual on-shell diagrams and describe the particular combinations that represent scattering amplitudes. We present a self-contained direct proof—using on-shell diagrams alone—that the BCFW construction of the all-loop integrand generates an object with precisely those singularities dictated by quantum field theory. We then show that the Grassmannian representation of loop integrands are always given in a remarkable "*d*log" form, which we illustrate using examples of simple one- and two-loop amplitudes. We discuss the implications of this representation for the transcendental functions that arise after the loop integrands are integrated.

We conclude our story in Chapter 18 with a discussion of a number of the outstanding open directions for further research.

2

Introduction to on-shell functions and diagrams

Theoretical explorations in field theory have been greatly advanced by focusing on interesting classes of observables—from local correlation functions and scattering amplitudes, to Wilson and 't Hooft loops, surface operators and line defects, to partition functions on various manifolds (see, e.g., [58, 59]). The central physical idea of our work is to extend the notion of "scattering amplitudes" to a broader class of objects called *on-shell functions*, which we introduce in this chapter.

2.1 On-shell particles, functions, and kinematical data

On-shell functions, like the S-matrix, depend only on the data describing physically observable external states. This data consists of the momentum $p_a^\mu \in \mathbb{R}^{3,1}$, mass m_a, spin σ_a, helicity $h_a \in \{\sigma_a, \sigma_a - 1, \ldots, -\sigma_a\}$, and any non-kinematical quantum numbers q_a that describe the external particle indexed by $a \in \{1, \ldots, n\}$. The momentum of any *observable* state satisfies the Einstein relation, $p^\mu p_\mu = m^2$; such particles are said to be "on the mass-shell" (the hyperboloid $p^\mu p_\mu - m^2 = 0$), or simply *on-shell*. In this work, we will focus on theories involving *massless* particles—those with $m_a = 0$; such particles can only have helicity $h_a = \pm\sigma_a$.

When an external particle is massless, the (2×2)-matrix constructed out of the Pauli matrices $\sigma_\mu^{\alpha\dot{\alpha}}$,

$$p_a^{\alpha\dot{\alpha}} \equiv p_a^\mu \sigma_\mu^{\alpha\dot{\alpha}} = \begin{pmatrix} p_a^0 + p_a^3 & p_a^1 - ip_a^2 \\ p_a^1 + ip_a^2 & p_a^0 - p_a^3 \end{pmatrix}, \tag{2.1}$$

(with entries labeled by $\alpha = 1, 2$, and $\dot{\alpha} = \dot{1}, \dot{2}$) will have a vanishing determinant:

$$\det(p_a^{\alpha\dot{\alpha}}) = (p_a^0)^2 - (p_a^1)^2 - (p_a^2)^2 - (p_a^3)^2 = 0 \; (\equiv \eta_{\mu\nu} p_a^\mu p_a^\nu); \tag{2.2}$$

and so $p_a^{\alpha\dot{\alpha}}$ has rank 1 (or 0). We can make this manifest via the substitution

$$p_a^{\alpha\dot{\alpha}} \equiv \lambda_a^\alpha \widetilde{\lambda}_a^{\dot{\alpha}} \Leftrightarrow \text{``}a\rangle[a\text{''}, \tag{2.3}$$

where $\lambda_a^{\alpha}, \widetilde{\lambda}_a^{\dot{\alpha}} \in \mathbb{C}^2$ are called *spinor-helicity* variables [60–64]. If the momentum p_a were *real*, we would have $\widetilde{\lambda}_a = \pm\lambda_a^*$; but we will often find it useful to allow all momenta to be complex, and consider $\lambda_a, \widetilde{\lambda}_a$ to be independent variables.

The rescaling $\lambda_a \mapsto t_a \lambda_a$, $\widetilde{\lambda}_a \mapsto t_a^{-1}\widetilde{\lambda}_a$ leaves the momentum p_a, (2.3), invariant and represents the action of the *little group* (for more details, see e.g. [2, 65]). Upon its complexification, the local Lorentz group becomes $SL(2)_L \times SL(2)_R$, with λ^{α} and $\widetilde{\lambda}^{\dot{\alpha}}$ transforming in the fundamental representations of $SL(2)_L$ and $SL(2)_R$, respectively. Knowing this, we may construct the Lorentz-invariants,

$$\langle a\, b\rangle \equiv \det\{\lambda_a, \lambda_b\} \quad \text{and} \quad [a\, b] \equiv \det\{\widetilde{\lambda}_a, \widetilde{\lambda}_b\}, \tag{2.4}$$

out of which all Lorentz-invariants of the momenta can be constructed—e.g.,

$$\eta_{\mu\nu}\, (p_a + p_b)^{\mu}\, (p_a + p_b)^{\nu} \equiv (p_a + p_b)^2 = \langle a\, b\rangle [a\, b], \tag{2.5}$$

for any on-shell, massless four-momenta p_a, p_b.

The Lorentz-invariant phase space ('LIPS') associated with an on-shell particle is given by the four degrees of freedom of $\lambda_a, \widetilde{\lambda}_a$ modulo the action of the little group—a $GL(1)$ redundancy. Thus, the differential form describing an on-shell particle's phase space can be written

$$d^3\text{LIPS}_a \equiv \frac{d^2\lambda_a\, d^2\widetilde{\lambda}_a}{\text{vol}(GL(1))}, \tag{2.6}$$

where "$1/\text{vol}(GL(1))$" represents the instruction to eliminate the $GL(1)$ redundancy of the little group, resulting in a three-dimensional form on phase space.

In general, we view all on-shell functions as being decorated with the Lorentz-invariant phase space measures for each external particle. As such, we may view them more formally as on-shell (differential) forms on the phase space of all the external kinematical data. Because of this, we will refer to on-shell functions interchangeably as on-shell forms throughout this work.

Let us denote the external wave function for particle a with helicity $h_a = \pm\sigma_a$ by '$|a\rangle^{h_a}$' (which should not be confused with the notation '$a\rangle$' for λ_a). Under the action of the little group, $\lambda_a \mapsto t_a\lambda_a, \widetilde{\lambda}_a \mapsto t_a^{-1}\widetilde{\lambda}_a$, the wave function transforms according to (see e.g. reference [2]):

$$|a\rangle^{h_a} \mapsto t_a^{-2h_a}|a\rangle^{h_a}. \tag{2.7}$$

Because of this, any Lorentz-invariant on-shell function of the external states (e.g. a scattering amplitude) must transform under the little group accordingly:

$$f(\ldots, t_a\lambda_a, t_a^{-1}\widetilde{\lambda}_a, h_a, \ldots) = t_a^{-2h_a} f(\ldots, \lambda_a, \widetilde{\lambda}_a, h_a, \ldots). \tag{2.8}$$

We will see in section 2.3 that this scaling property together with momentum conservation uniquely fixes the kinematical dependence of the S-matrix for three massless particles with any helicities (to all loop orders!).

Before moving on to the main subject of this chapter, on-shell diagrams, it will be helpful to establish some useful notational conventions. As we have seen, the kinematical data describing on-shell particles are specified by a pair of two-vectors $\lambda_a, \widetilde{\lambda}_a \in \mathbb{C}^2$ for each particle $a \in \{1,\ldots,n\}$. We will frequently find it convenient to organize these data collectively into a pair of $(2\times n)$ matrices written according to the following convention:

$$\{\lambda_a^\alpha\} \Leftrightarrow \lambda \equiv \begin{pmatrix} \lambda_1^1 \cdots \lambda_n^1 \\ \lambda_1^2 \cdots \lambda_n^2 \end{pmatrix} \equiv \begin{pmatrix} \lambda^1 \\ \lambda^2 \end{pmatrix} \equiv \big(\lambda_1 \cdots \lambda_n\big), \tag{2.9}$$

and similarly for $\widetilde{\lambda}$. That is, 'λ' collectively denotes all external spinors and components $\{\lambda_a^\alpha\}$, 'λ_a' denotes the two-vector of spinor components for the ath particle, and 'λ^α' denotes the n-vector consisting of each spinor's α-component.

Finally, momentum conservation must always be satisfied for the particles involved in any scattering process. If we conventionally take all the external momenta to be incoming, then momentum conservation becomes the constraint

$$\delta^{2\times 2}\left(\sum_{a=1}^{n} \lambda_a^\alpha \widetilde{\lambda}_a^{\dot{\alpha}}\right) \equiv \delta^{2\times 2}\big(\lambda\cdot\widetilde{\lambda}\big), \tag{2.10}$$

where we have introduced "$\cdot$" to denote a summation over external particles. This will prove a useful convention throughout the rest of this work. Notice also that because the four-momenta are each written as a (2×2) matrix according to (2.1), momentum conservation is naturally organized into a (2×2) matrix of constraints. Whenever a system of constraints is naturally organized into a $(k\times m)$ matrix, we will indicate this by writing "$\delta^{k\times m}\big(\cdots\big)$" as in (2.10).

We should clarify that, throughout this work, we consider only "holomorphic" δ-functions. These behave very similarly to the traditional "δ-functions" of quantum mechanics, but require no use of a complex norm. Concretely, our δ-functions are defined as *residue* (or *contour*) prescriptions according to:

$$\int dz\, g(z)\, \delta\big(f(z)\big) \equiv \sum_{z^*|f(z^*)=0} \mathrm{Res}\left(\frac{g(z)}{f(z)}; z^*\right). \tag{2.11}$$

The following is a simple, illustrative example:

$$\int dz\, \delta\big(f(z)\big) = \sum_{z^*|f(z^*)=0} \frac{1}{f'(z^*)}, \tag{2.12}$$

where $f'(z^*)$ denotes the derivative of f with respect to z, evaluated at z^*. The reader should notice that δ-functions defined in this way behave the same way as ordinary δ-functions, except that, in (2.12) for example, no absolute-value sign appears around the derivative of f. This definition generalizes to the case $\delta^{k\times m}\big(\cdots\big)$, which specifies the contour for a co-dimension $(k\times m)$ residue (see e.g. [66]).

Thus, if we re-write the four components of the δ-functions encoding momentum conservation, (2.10), in terms of Lorentz-invariant quantities by contracting with some reference null momenta $\{\lambda_r\widetilde{\lambda}_r, \lambda_s\widetilde{\lambda}_s\}$, we pick-up an ordinary *Jacobian* from the change of variables (without any absolute-value signs!):

$$\delta^{2\times2}(\lambda\cdot\widetilde{\lambda}) = \langle rs\rangle[rs]\,\delta\Big(\sum_a\langle ra\rangle[ar]\Big)\delta\Big(\sum_a\langle ra\rangle[as]\Big)\delta\Big(\sum_a\langle sa\rangle[ar]\Big)\delta\Big(\sum_a\langle sa\rangle[as]\Big). \tag{2.13}$$

2.2 Scattering amplitudes and their amalgamations

We would like to understand the entire class of physically meaningful functions describable *exclusively* in terms of on-shell kinematical data for some number of external states (without any reference to virtual particles, gauge redundancies, ghosts, or any of the other unphysical baggage associated with the traditional approach to quantum field theory).

A particularly important example of this class of functions is the full scattering amplitude (the "S-matrix") $\mathcal{A}_n(\lambda,\widetilde{\lambda},h)$ for n particles with momenta given by $\lambda,\widetilde{\lambda}$ and with helicities $h\equiv(h_1\cdots h_n)$. But scattering amplitudes represent only a small subset of the meaningful gauge-invariant functions that can be constructed exclusively in terms of on-shell external data. In particular, knowing even a few scattering amplitudes, we can "amalgamate" them into more complex objects we call *on-shell functions*, which can be represented as *on-shell diagrams*, constructed out of amplitudes by sewing them together in a natural way.

The most familiar example of an on-shell function built out of scattering amplitudes, but which is not a scattering amplitude itself, is known as a *factorization channel*; diagrammatically, we represent a factorization channel as follows

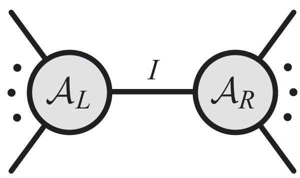

(2.14)

A factorization channel is well defined without any reference to any off-shell degrees of freedom, and corresponds to the particular on-shell function

$$\sum_{h_I}\int\frac{d^2\lambda_I\,d^2\widetilde{\lambda}_I}{\mathrm{vol}(GL(1))}\mathcal{A}_L(\ldots,\lambda_I,\widetilde{\lambda}_I,h_I)\mathcal{A}_R(\lambda_I,-\widetilde{\lambda}_I,-h_I,\ldots). \tag{2.15}$$

(Here, we have left implicit a summation over any non-kinematical quantum numbers labeling the internal particle.) Notice that in (2.15), integration over the internal particle's on-shell phase space is trivial: it is *completely localized* by the momentum-conserving δ-functions present in the two amplitudes. Notice, however, that even after localizing the three-dimensional phase space integral, five of the eight initial δ-functions remain; these impose the four constraints of overall momentum conservation, together with the additional constraint $\big(\sum_{a\in L}p_a\big)^2=0$.

We can understand how the formula (2.15) follows from locality and unitarity as follows. Because the internal particle is massless (and on-shell), the two scattering amplitudes can involve particles separated by arbitrary distances in space and time; as such, locality dictates that they must be independent, and so the amplitudes must be multiplied together. And unitarity instructs us to marginalize over any unobserved states—integrating over each internal particle's on-shell phase space, (2.6), and summing over possible helicities (and any other quantum numbers).

This rule can be generalized to define arbitrary graphs built out of on-shell scattering amplitudes separated by internal (but on-shell) particles. Thus, the on-shell function associated with a graph Γ involving an amplitude $\mathcal{A}_v$ at each vertex v and any number of internal particles $i \in I$ is defined according to

$$f_\Gamma \equiv \prod_{i\in I}\left(\sum_{h_i}\int\frac{d^2\lambda_i\, d^2\widetilde{\lambda}_i}{\mathrm{vol}(GL(1))}\right)\prod_v \mathcal{A}_v, \tag{2.16}$$

where convention that amplitudes involve only *incoming* momenta dictates that two ends of any internal line must involve opposite momenta and helicities.

One important physical characteristic of any on-shell diagram obtained in the way described above is the number n_δ of constraints (if any) that are imposed on the external kinematical data beyond overall momentum conservation. Because each vertex amplitude imposes momentum conservation, we start with a total $4n_V$ constraints for a diagram with n_V vertices; these constraints *always* imply overall momentum conservation, and so $4n_V - 4$ of the constraints can contribute to n_δ. But we must integrate over each of the n_I internal particles' three dimensional on-shell phase space; therefore, if as many of these phase space integrals can be localized as possible, then the number of excess constraints would be given by

$$n_\delta \equiv 4n_V - 3n_I - 4. \tag{2.17}$$

When $n_\delta = 0$, it is possible that all the phase space integrals can be localized by the δ-functions, resulting in an on-shell diagram that neither imposes any constraints on the external kinematics, nor leaves us with any remaining integrals over phase space to perform. When this is the case, the on-shell diagram is simply a (rational) function of the external kinematical variables; such on-shell functions have historically been called "leading singularities" in the physics literature [12,67], for reasons we will discuss momentarily. When $n_\delta > 0$, however, the resulting on-shell function imposes some number of excess constraints on the external kinematics; these objects have historically been called "singularities" or said to have "singular support". Finally, when $n_\delta < 0$, there are internal phase space integrations that *cannot* be localized by the δ-functions, requiring that we specify a choice of contour for their integration. We will see in section 2.6 that with the appropriate choice of contours, these on-shell phase space integrations *exactly* reproduce the "loop-level" contributions in the Feynman expansion (where, in a

Feynman diagram, this integration would be performed over degrees of freedom associated with the *off-shell* "loop momenta" of *virtual* particles).

We have already seen an example of a diagram for which $n_\delta > 0$: the factorization channel, (2.14). The following is an example of a diagram with $n_\delta = 0$:

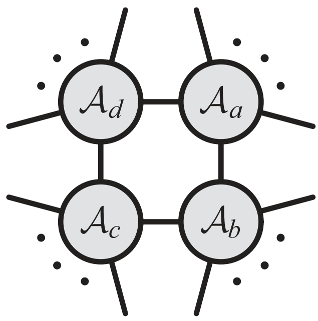

(2.18)

For this diagram, *all* the internal phase space integrations are localized by momentum conservation at the vertices; and so, (2.18) involves no internal degrees of freedom and imposes no additional constraints on the external kinematical data. Because of this, the diagram (2.18) represents an ordinary (rational) function of the external momenta (and has nothing directly to do with a "loop" in the sense of traditional, off-shell perturbation theory).

As mentioned earlier, the diagram (2.18) would have historically been called a "one-loop leading singularity"; this is because it can be interpreted as a co-dimension–four residue of the one-loop Feynman integrand—where each residue constrains an *off-shell* internal (hence virtual) particle to be put *on-shell*. We choose not to use this terminology here, because it subordinates on-shell functions like (2.18) relative to objects such as Feynman diagrams defined in terms of unphysical, off-shell degrees of freedom. As emphasized above, *all* on-shell diagrams are intrinsically well defined, without any reference to Feynman diagrams.

2.3 On-shell building blocks: the massless three-particle *S*-matrix

As we saw in section 2.2, on-shell scattering amplitudes can be "amalgamated" into more complicated on-shell objects; but we cannot begin to investigate such objects until we know some scattering amplitudes to feed into the machinery. In this section, we show that the *S*-matrix involving three massless particles (with arbitrary helicities) is uniquely fixed by first principles *to all orders of perturbation theory*. Given just these basic amplitudes as building blocks, the procedure described above immediately leads to an incredible variety of on-shell diagrams and corresponding functions, defined to all orders of perturbation theory. Moreover, as we will see in section 2.6, the entire perturbative *S*-matrix (at least in