

THE ENUMERATION
OF PLANE
TROPICAL CURVES

Hannah Markwig

Beim Fachbereich Mathematik
der Technischen Universität Kaiserslautern
zur Verleihung des akademischen Grades
Doktor der Naturwissenschaften
(Doctor rerum naturalium, Dr. rer. nat.)
eingereichte Dissertation.

1. Gutachter: JProf. Dr. Andreas Gathmann
2. Gutachter: Prof. Dr. Bernd Sturmfels

D 386

FRAGE AN HANNAH:

Wieviele Kurven sind das?



eine

fünf

vier

sechs



CONTENTS

1. Introduction	1
1.1. Tropical geometry	1
1.2. Enumerative geometry	2
1.3. Tropical enumerative geometry	3
1.4. The content of this thesis	5
1.5. Acknowledgments	5
2. Motivation on plane tropical curves	7
2.1. The field of Puiseux series	8
2.2. The tropical semiring	10
2.3. A combinatorial description of tropical curves	13
3. Classical concepts for the enumeration of plane curves	17
3.1. The moduli space of stable maps	18
3.2. Kontsevich's formula to determine $N_{\text{cplx}}^{\text{irr}}(d, 0)$	24
3.3. The algorithm of Caporaso and Harris to determine $N_{\text{cplx}}(d, g)$	31
3.4. A short overview of curves on toric surfaces	42
4. The tropical enumerative problem in the plane	47
4.1. Abstract tropical curves	47
4.2. Parametrized tropical curves	51
4.3. The comparison of parametrized tropical curves with the definitions of tropical curves from chapter 2	53
4.4. The moduli space of parametrized tropical curves	57
4.5. The tropical enumerative problem	66
4.6. The moduli space and the structure of polyhedral complexes	71
4.7. The proof of theorem 4.53	75
5. The correspondence of tropical curves and lattice paths	90
5.1. Properties of parametrized tropical curves described in the dual language of Newton subdivisions	91
5.2. The dual of a marked parametrized tropical curve	99
5.3. Lattice paths	105

5.4.	The correspondence of tropical curves and lattice paths	107
6.	The correspondence of complex curves and tropical curves	114
6.1.	J_t -holomorphic curves	115
6.2.	The “limit” of an amoeba — amoebas of J_t -holomorphic curves	117
6.3.	The number of complex curves whose limit is a given tropical curve	121
7.	A tropical proof of Kontsevich’s formula	129
7.1.	The enumerative problem for rational parametrized tropical curves	130
7.2.	Tropical forgetful maps	132
7.3.	Reducible curves and Kontsevich’s formula	138
8.	The tropical Caporaso-Harris algorithm	150
8.1.	Tropical curves that satisfy higher order tangency conditions to a line	151
8.2.	The tropical Caporaso-Harris algorithm	153
8.3.	The tropical Caporaso-Harris algorithm for irreducible curves	157
8.4.	The correspondence of complex curves tangent (of higher order) to a line and tropical curves with ends of higher weight	159
9.	The Caporaso-Harris algorithm in the lattice path setting	165
9.1.	Generalized lattice paths	165
9.2.	The Caporaso-Harris algorithm for generalized lattice paths	167
9.3.	The correspondence between tropical curves with ends of higher weight and generalized lattice paths	173
	References	175
	Index	177

1. INTRODUCTION

Tropical geometry is a rather new field of algebraic geometry. The main idea is to replace algebraic varieties by certain piece-wise linear objects in \mathbb{R}^n , which can be studied with the aid of combinatorics. There is hope that many algebraically difficult operations become easier in the tropical setting, as the structure of the objects seems to be simpler.

In particular, tropical geometry shows promise for application in enumerative geometry. Enumerative geometry deals with the counting of geometric objects that are determined by certain incidence conditions. Until around 1990, not many enumerative questions had been answered and there was not much prospect of solving more. But then Kontsevich introduced the moduli space of stable maps which turned out to be a very useful concept for the study of enumerative geometry. The idea of Kontsevich was motivated by physics, more precisely, by string theory. Since then, enumerative geometry has gained a lot more attention: not only from physicists, but also from mathematicians, as the theory of stable maps has become rich and elaborated. However, a lot of questions remain open, and there are still many mathematicians working in enumerative geometry.

Tropical geometry supplies many new ideas and concepts that could be helpful to answer enumerative problems. However, as a rather new field, tropical geometry has to be studied more thoroughly. This thesis is concerned with the “translation” of well-known facts of enumerative geometry to tropical geometry. We will first give a short introduction to tropical geometry and then explain the well-known results of enumerative geometry that will be “translated” in this thesis.

1.1. TROPICAL GEOMETRY

Tropical geometry is so far best developed for plane curves. An idea Kontsevich proposed and Mikhalkin elaborated in [23] is to apply the map

$$\text{Log} : (\mathbb{C}^*)^2 \rightarrow \mathbb{R}^2 : (z, w) \mapsto (\log |z|, \log |w|)$$

to a complex curve in a toric surface. The observation is that the image of a complex curve under this map looks roughly like a graph in \mathbb{R}^2 with linear edges. When we shrink the image to a certain limit, we end up with such a graph fulfilling a condition called the “balancing condition”. Such a graph will be referred to as a tropical curve. An analogous “deformation” of the complex numbers yields a semiring $(\mathbb{R} \cup \{-\infty\}, \max, +)$ with operations \max as addition and $+$ as multiplication. This semiring has been known to computer scientists before and is referred to as “tropical semiring” in honour of the Brazilian mathematician and computer scientist Imre Simon (see for example [29]).

As already mentioned, tropical curves look, shortly described, like graphs in \mathbb{R}^2 which fulfill certain conditions. The balancing condition allows to associate a dual to a tropical curve, which is a regular subdivision of a lattice polygon in \mathbb{Z}^2 (see section 2.3). Therefore, the data of a tropical curve can be described purely combinatorially using lattice polygons (respectively, their dual graphs).

A lot of work has been done to “translate” classical concepts to this tropical setting, especially in enumerative geometry. A well-known problem of enumerative geometry is to determine the numbers $N_{\text{cplx}}(d, g)$ of complex genus g plane curves of degree d passing

through $3d+g-1$ points in general position. In [23], Mikhalkin associates a multiplicity to each tropical curve C , which coincides with the number of complex curves that project to C (under Log and taking the limit). He shows that the number $N_{\text{cplx}}(d, g)$ is equal to the number $N_{\text{trop}}(d, g)$ of tropical curves through $3d+g-1$ points, counted with multiplicity. This important result is referred to as Correspondence Theorem. We will describe it more precisely in chapter 6. Furthermore, he computes $N_{\text{cplx}}(d, g) = N_{\text{trop}}(d, g)$ purely combinatorially using certain lattice paths in the lattice polygon dual to the tropical curves (see chapter 5). Siebert and Nishinou extended the Correspondence Theorem to rational curves in an n -dimensional toric variety [24]. Shustin showed the same for certain singular plane curves [26].

But not only enumerative geometry has been translated to the tropical world: Izhakian found an analogue to the duality of curves [16], Vigeland established a group law on tropical elliptic curves [33] and Tabera dealt with a tropical Pappus' Theorem [30], just to mention a few.

Tropical research is not restricted to the translation of classically well-known facts, there are actually new results shown by means of tropical geometry that have not been known before. For example, Mikhalkin gave a tropical algorithm to compute the Welschinger invariant for real curves [23] and Itenberg, Shustin and Kharlamov were able to estimate the Welschinger invariant (even for large degrees where the computation using Mikhalkin's algorithm is too complicated) using tropical curves [15]. Furthermore, there are ideas by Mikhalkin to compute Zeuthen numbers, that is the numbers of plane curves which do not only satisfy the condition to pass through certain points, but also tangency conditions to lines.

This shows that tropical geometry can indeed be a tool for a better understanding of classical geometry.

1.2. ENUMERATIVE GEOMETRY

As already mentioned, enumerative geometry deals with the counting of geometric objects that satisfy given incidence conditions. The conditions must be chosen in such a way that there is actually a finite number of objects that satisfy them.

The main strategy to count objects is to construct a moduli space which parametrizes these objects. The special objects that satisfy one of the given incidence conditions will then correspond to a subspace of the moduli space. In order to count objects that satisfy all conditions, we have to intersect the subspaces corresponding to each condition, and determine the number of points in the 0-dimensional intersection product. (The way the incidence conditions were chosen - such that there are only finitely many objects satisfying them - guarantees that this intersection is indeed 0-dimensional.) The moduli space of stable maps is a moduli space that seems appropriate for a large class of enumerative questions. Therefore, enumerative geometry deals basically with intersection theory on this moduli space. The numbers which occur as intersection numbers on the moduli space of stable maps are called Gromov-Witten invariants. The numbers $N_{\text{cplx}}(d, g)$ are called Gromov-Witten invariants of \mathbb{P}^2 , because they arise as intersection numbers on the moduli space of stable maps to \mathbb{P}^2 . For genus 0, they were first computed by Kontsevich using

the moduli space of rational stable maps. This moduli space is based on a different understanding of “curve”. A curve can on the one hand be considered as an embedded object in the ambient space. On the other hand, it can be considered as a map from an abstract curve to the ambient space. The moduli space of stable maps parametrizes such maps (satisfying certain conditions which we do not want to make precise here — the definition can be found in section 3.1). The boundary of this moduli space consists of maps where the underlying abstract curve is reducible. Kontsevich shows that two special divisors which are contained in the boundary of the moduli space are linearly equivalent. We can intersect both divisors with some more conditions to get something zero-dimensional. The stable maps which are contained in these two zero-dimensional subsets are reducible. Comparing the image curves on both sides yields a recursive formula that determines $N_{\text{cplx}}(d, 0)$ depending on the numbers $N_{\text{cplx}}(d', 0)$ for $d' < d$. The idea how Kontsevich’s formula can be derived using the moduli space of stable maps is described more precisely in section 3.2. Note that due to the special structure of the moduli space of rational abstract curves the methods of this proof cannot be generalized to higher genus.

For arbitrary genus, there is an algorithm developed by Caporaso and Harris that determines $N_{\text{cplx}}(d, g)$ using degenerations of curves after specializing the points [4]. The algorithm does not only involve the numbers $N_{\text{cplx}}(d, g)$, but also the numbers of curves that do not only pass through the appropriate number of points in general position, but in addition satisfy tangency conditions (of higher order) to a fixed line L . These numbers are referred to as “relative Gromov-Witten invariants”. By specializing the points one by one to lie on the line L , Caporaso and Harris derive recursive relations between the relative Gromov-Witten invariants, that allow to compute $N_{\text{cplx}}(d, g)$. For example, if already d points are specialized to lie on L and we move the $d + 1$ -st point to L , the curves passing through this configuration of points can no longer be irreducible. Instead, L has to split off as a component of the curve. The remaining component then is a curve of degree $d - 1$, which may fulfill tangency conditions of arbitrary order to L . The recursive formula sums up the possibilities for a curve passing through a given set of points after we moved one point to L : there are irreducible curves which then pass through one more point on L , and there are the above described reducible curves. The difficult part of the algorithm is to determine which reducible curves appear and with which multiplicity they contribute (as components of a class in the Chow group of the moduli space). A more detailed description of the ideas of the Caporaso-Harris algorithm - though not a proof - can be found in section 3.3.

The aim of this thesis is to translate both Kontsevich’s formula and the Caporaso-Harris algorithm to tropical geometry.

1.3. TROPICAL ENUMERATIVE GEOMETRY

The Correspondence Theorem of Mikhalkin mentioned above shows that the numbers $N_{\text{trop}}(d, g)$ coincide with the numbers $N_{\text{cplx}}(d, g)$. Furthermore, we have seen that the numbers $N_{\text{cplx}}(d, g)$ satisfy some recursive relations. Of course, the Correspondence Theorem shows that the numbers $N_{\text{trop}}(d, g)$ satisfy the same relations. The aim of this thesis is to reprove this fact without using Mikhalkin’s Correspondence Theorem - that is, to find “tropical proofs” of Kontsevich’s formula and the Caporaso-Harris algorithm. We will

shortly describe the questions and challenges that arise in this context.

A fact which is classically well-known is that the number of complex curves through the appropriate number of points does not depend on the position of the points, as long as it is sufficiently general. Again, the Correspondence Theorem shows immediately that the analogue is true for the numbers $N_{\text{trop}}(d, g)$. However, the statement was not shown purely with methods from tropical geometry so far. When trying to prove a tropical Caporaso-Harris algorithm, we need to specialize the points. Therefore it is important to see that the number of tropical curves through $3d + g - 1$ points does not depend on the position of the points (see theorem 4.53). We prove this statement within tropical geometry in section 4.7, using the moduli space of tropical curves. The definition of the moduli space of tropical curves we use here is inspired by ideas of Mikhalkin. We define it in section 4.2 and discuss it further in section 4.4. The proof is more complicated than the proof that the numbers $N_{\text{cplx}}(d, g)$ do not depend on the position of the points, because we have to count the tropical curves with multiplicity as mentioned above. If we count each curve without that factor, this number will in fact depend on the position of the points.

With the aid of the moduli space of tropical curves and using analogous ideas as for the classical proof of Kontsevich's formula, we give a tropical proof that the numbers $N_{\text{trop}}(d, 0)$ satisfy Kontsevich's formula. This result is described in section 7. We do not define tropical analogues of divisors and intersection theory. In our tropical proof, other methods replace these concepts. We hope that our methods will help to understand tropical geometry in a more general context and maybe to define tropical divisors or intersection theory later.

For the Caporaso-Harris algorithm, we do not have to work with a tropical moduli space. To the contrary, it is sufficient to consider a tropical curve through a given set of points and to study the possible degenerations that arise after specializing the points. However, as mentioned above, the algorithm involves relative Gromov-Witten invariants. Therefore we have to define tropical analogues of relative Gromov-Witten invariants first (see section 8.1). The tropical proof of Caporaso's and Harris' algorithm can be found in chapter 8. Mikhalkin's Correspondence Theorem states that the numbers of complex and of tropical curves through a given set of points coincides, but it does not state an analogue for the relative Gromov-Witten invariants. The equality of the numbers of complex and tropical curves passing through points and satisfying additional tangency conditions can be derived recursively as we know that both numbers fulfill the relations given by Caporaso and Harris — starting with the fact that there is one complex as well as one tropical line through 2 points. However, we wish to give a more direct proof of this correspondence, too. This result can be found in section 8.4.

As mentioned before, tropical curves are dual to lattice polygons. Even more, Mikhalkin gave a way to count tropical curves using certain paths in the lattice polygon dual to the curve. Of course, the numbers of these paths fulfill the recursive relations known from classical geometry, too. But again, we were interested in finding a direct proof of these recursions. As before, we first had to define analogues of relative Gromov-Witten invariants for lattice paths, and then reprove the Caporaso-Harris formula in the lattice path setting. Also, we can see directly that our generalized lattice paths correspond to tropical curves satisfying tangency conditions to a line. These results can be found in chapter 9.

To sum up, we succeeded in “translating” the concepts from enumerative geometry mentioned above to tropical geometry. We believe that our work is a step towards a better understanding of tropical geometry, which will in our opinion result in a better understanding of classical enumerative geometry.

1.4. THE CONTENT OF THIS THESIS

This thesis is organized as follows. In chapter 2, we motivate what tropical curves should look like and give an overview about some of their basic properties that were known before. In chapter 3, we give a short introduction to enumerative geometry, and, more precisely, to the concepts we want to translate to the tropical world later on. In chapter 4, we define tropical curves combinatorially and introduce the moduli space of tropical curves. We use it to derive the fact that the numbers $N_{\text{trop}}(d, g)$ do not depend on the position of the points without passing to the complex world using the Correspondence Theorem. In chapter 5, we give an overview of the duality of tropical curves and lattice paths found by Mikhalkin, because we want to generalize these ideas to relative Gromov-Witten invariants in chapter 9. In chapter 6, we state Mikhalkin’s Correspondence Theorem and give a short overview of the structure of his proof. In chapter 7, we reprove Kontsevich’s formula tropically. In chapter 8, we proceed towards the Caporaso-Harris algorithm in the tropical setting. We define tropical relative Gromov-Witten invariants, prove the Caporaso-Harris algorithm tropically, and prove a generalized Correspondence Theorem for relative Gromov-Witten invariants. In chapter 9 finally, we reprove the Caporaso-Harris algorithm in the lattice path setting.

That is, the new results can be found in the chapters 4, 7, 8 and 9. In the other chapters known facts are described that we need for our work.

All results mentioned above were derived in joint work with my advisor Andreas Gathmann. They were published (some so far only as preprint) in [12], [13] and [14]. In this joint work it is not easy to separate the contributions we both made. As far as it can be told, most important ideas of 4.4 were developed by Andreas Gathmann, of 4.7 by myself. (The main results of chapter 4 can be found in 4.4 and 4.7.) In chapter 7, important ideas of both of us were used. The basic ideas of chapter 8.1 and 8.2 are due to Andreas Gathmann, whereas the ideas of 8.3 and 8.4 are due to myself. The main ideas of chapter 9 are due to myself.

1.5. ACKNOWLEDGMENTS

First of all I would like to thank my advisor Andreas Gathmann. Andreas was a really great advisor from whom I could learn a lot — about mathematics, but also about the way to teach mathematics. Many thanks to Andreas for his support, for many helpful discussions, for providing me with ideas, for correcting my mistakes, for giving me the opportunity to attend many interesting conferences ...

Andreas also managed to establish a whole working group of tropical geometry in Kaiserslautern. I owe thank to this group of students, too. Many problems were solved after discussions with Christian Eder, Michael Kerber and Johannes Rau. Thanks to Eric

Westenberger who introduced me to tropical geometry. I also gained a lot of understanding after discussions with Grisha Mikhalkin, Bernd Sturmfels and Oliver Wienand. Many thanks to Ilya Itenberg who pointed out a serious mistake in an earlier version of the proof of the Caporaso-Harris algorithm for lattice paths.

Many friends invested a lot of time in proofreading. As I know that proofreading is not really an amusement, I am particularly thankful to all who still did that for me: many thanks to Michael Brickenstein, Michael Cuntz, Christian Eder, Andreas Gathmann, Michael Kerber, Thomas Markwig and Johannes Rau.

Furthermore, I want to thank my family and friends for their support and encouragement. In particular, I would like to mention Michael Brickenstein, who provided me with an advertising shirt for tropical geometry; Nadine Cremer, a “fellow sufferer” in the last days; Michael Kunte, whose patience to listen is unbelievable; and Ingo Münch, who often surprised me with nice interpretations of my black board writings. Furthermore, I want to thank my parents Uschi and Micha Markwig for their selfless support over many years. Finally, I want to thank my beloved husband Thomas for supporting and enduring me, for encouraging me even when I went through a fainthearted period, for patient help with computer questions, for a lot of help with the design of the title page ... — for sharing his life with me.


To conclude, I would like to thank whoever invented the term “tropical”, because I think that this term reflects just perfect how much fun it is.

2. MOTIVATION ON PLANE TROPICAL CURVES


The idea what a plane tropical curve should be arises after applying the map

$$\text{Log} : (\mathbb{C}^*)^2 \rightarrow \mathbb{R}^2 : (s, t) \mapsto (\log |s|, \log |t|)$$

to a complex curve in $(\mathbb{C}^*)^2$. As a first example, let L be a projective line in $\mathbb{P}_{\mathbb{C}}^2$, and apply Log to the restriction of L to $(\mathbb{C}^*)^2$. Let $(x : y : z)$ be the coordinates of \mathbb{P}^2 , and identify \mathbb{C}^2 with the set $\{z \neq 0\}$. Then the map Log associates the point $(\log |\frac{x}{z}|, \log |\frac{y}{z}|) \in \mathbb{R}^2$ to a point $(x : y : z) \in \mathbb{P}^2$. The line L intersects the coordinate line $\{x = 0\}$ in one point. When we move along the line L towards the intersection with $\{x = 0\}$, the first coordinate of the image point under Log will tend to $-\infty$. Also, when we move towards the intersection with the coordinate line $\{y = 0\}$, the second coordinate of the image will tend to $-\infty$. When we move towards the intersection with $\{z = 0\}$, both coordinates will become big and their difference will tend to a constant. Furthermore, the image $\text{Log}(L) \subset \mathbb{R}^2$ should be something 2-dimensional, as the complex line has two real dimensions. These observations suggest that the image will look similar to the following:




The image $\text{Log}(L)$ is called the amoeba of the line L . A tropical line can be thought of as a limit of this amoeba. In fact, the tropical line is what we get after shrinking the amoeba to something one-dimensional. It looks like the amoeba from very, very far away:




The only information kept are the three infinite rays and their directions.

2.1 Remark

Note that the primitive integer vectors pointing in these three directions sum up to 0:



As a second example, let $C \subset \mathbb{P}^2$ be a conic. As before, let us examine what happens at the coordinate lines. C intersects $\{x = 0\}$ in two points, $(0 : p_0 : 1)$ and $(0 : p_1 : 1)$. So we can move along C near p_0 and the first coordinate of the image will tend to $-\infty$, whereas the second tends to $\log |p_0|$. When we move along C near p_1 , the first coordinate will again tend to $-\infty$, but the second to $\log |p_1|$. With the same argument as above, we can see that the amoeba of a conic will have two “tentacles” in each of the three directions $(-1, 0)$, $(0, -1)$ and $(1, 1)$. However, we can not say precisely what happens in the middle. When we try again to shrink the amoeba to something 1-dimensional to get an idea of how a tropical conic should look like, there are indeed several possibilities of what can happen in the middle.



The picture shows three different types of a tropical conic.

The two examples of a tropical line and a tropical conic suggest that a tropical curve should be a piece-wise linear object which is in some sense the image of a complex curve — under the map Log and a degeneration process which will be specified in chapter 6. We can hope now that tropical curves still carry a lot of properties that the original complex curves had, and furthermore, that they are easier to deal with due to their linearity. In fact, we will see in this thesis that a lot of concepts of classical curves can be “translated” to tropical geometry, and that proofs are in general easier in the tropical setting.

The idea that a tropical curve can be thought of as a limit of an amoeba serves as a motivation why tropical curves are interesting and what they should roughly look like. However, we will not use it as a definition, as it is rather difficult to make the notion of limit precise. Instead, we will use a different approach to define tropical curves and see in chapter 6 that the objects of our definition are actually limits of amoebas of complex curves. (The notion of limit will be made precise in chapter 6).

2.1. THE FIELD OF PUISEUX SERIES

An idea which leads us to a possible definition of a tropical curve is to replace the field \mathbb{C} by another algebraically closed field K (of characteristic 0). As the main properties of curves should not depend on the algebraically closed field we are working with, we can as well consider curves over K and hope to define tropical curves somehow as an image of a curve over K . What is important about our choice of field is that we have again a map from K to \mathbb{R} which is in some sense similar to \log . However, the field should in another

sense be even “better” than \mathbb{C} : it shall be equipped with a norm which is non-archimedean (which does not hold for the canonical norm on \mathbb{C}). We will specify the requirements we have on the field K and the special map similar to \log .

2.2 Definition

Let K be any field. A map $|\cdot| : K \rightarrow \mathbb{R}_{\geq 0}$ satisfying

- $|a| = 0$ if and only if $a = 0$,
- $|ab| = |a| \cdot |b|$ and
- $|a + b| \leq \max\{|a|, |b|\}$ for all $a, b \in K$

is called a *non-archimedean norm*.

Note that if $L|K$ is an extension of a finite degree, the norm extends uniquely to a norm on L . Especially, as every element a in the algebraic closure of K is contained in an extension of a finite degree over K , we can extend the norm to a . Altogether, we get an extension of the norm to the algebraic closure of K (see [6], chapter 1.)

2.3 Definition

A *valuation* is a map $\text{val} : K \rightarrow \mathbb{R} \cup \{-\infty\}$ satisfying

- $\text{val}(a) = -\infty$ if and only if $a = 0$,
- $\text{val}(ab) = \text{val}(a) + \text{val}(b)$ and
- $\text{val}(a + b) \leq \max\{\text{val}(a), \text{val}(b)\}$ for all $a, b \in K$.

Non-archimedean norms are in bijection with valuations by $\text{val}(a) = \log |a|$.

Now we ask K to be a complete algebraically closed non-archimedean field. That is, K is algebraically closed and there is a valuation $\text{val} : K \rightarrow \mathbb{R} \cup \{-\infty\}$ such that e^{val} defines a norm on K . Furthermore, K has to be complete with respect to the norm e^{val} .

We can then as before consider curves C in $(K^*)^2$ and their image in \mathbb{R}^2 under the map

$$\text{Val} : (K^*)^2 \rightarrow \mathbb{R}^2 : (a, b) \mapsto (\log |a|, \log |b|) = (\text{val}(a), \text{val}(b)).$$

The main example of such a field K is the *completion of the field of Puiseux series*. Take the algebraic closure $\overline{\mathbb{C}((t))}$ of the field $\mathbb{C}((t))$ of Laurent series. An element of $\overline{\mathbb{C}((t))}$ is a Puiseux series

$$p(t) = a_1 t^{q_1} + a_2 t^{q_2} + a_3 t^{q_3} + \dots$$

where $a_i \in \mathbb{C}$ and $q_1 < q_2 < q_3 < \dots$ are rational numbers with bounded denominators. Set $\text{val}(p(t)) = -q_1$. Now define K to be the completion of $\overline{\mathbb{C}((t))}$ with respect to the norm e^{val} . The valuation extends to this completion by $\text{val} = \log |\cdot|$.

2.4 Definition

Let $C \subset (K^*)^2$ be a curve, where K is the completion of the field of Puiseux series, as above. Then define the *tropical curve associated to C* as the closure of the image $\text{Val}(C) \subset \mathbb{R}^2$ of C .

Note that for complex curves, we wanted to define tropical curves as a limit of the amoeba, that is, a limit of the image $\text{Log}(C)$. Unlike that, our definition 2.4 does not include any limit, but only the closure of the image $\text{Val}(C) \subset \mathbb{R}^2$. The reason is that $\overline{\text{Val}(C)}$ differs

from the amoeba of a complex curve. It is not a 2-dimensional object with tentacles, but it does in fact look like the limit of an amoeba: it is a 1-dimensional polyhedral complex. To see this, we examine what the map val does to the field structure $(K, +, \cdot)$ and define tropical curves algebraically as something analogous to a zeroset of an ideal of “ $\text{val}(K, +, \cdot)$ ”.

2.2. THE TROPICAL SEMIRING

The aim of this section is to describe tropical curves algebraically, similar to a zeroset of a polynomial. The observation what the map val does to the field $(K, +, \cdot)$ will help to find out which ground field (resp. ground structure, as we will not end up with a field, actually) will be appropriate for this aim. First of all, the image of the map val is of course $\mathbb{R} \cup \{-\infty\}$. Definition 2.3 gives us an idea what happens to the operations “+” and “ \cdot ”: they become “max” and “+”. We will therefore work with the following ground structure:

2.5 Definition

The *tropical semiring* $(\mathbb{R} \cup \{-\infty\}, \oplus, \odot)$ is the semiring with underlying set $\mathbb{R} \cup \{-\infty\}$ and operations *tropical addition* and *tropical multiplication* defined by

$$x \oplus y := \max\{x, y\} \quad \text{and} \quad x \odot y := x + y \quad (x, y \in \mathbb{R} \cup \{-\infty\}).$$

The addition is idempotent in the sense that $a \oplus a = a$. The extension of \mathbb{R} by $-\infty$ provides us with a neutral element of addition. However, there are no inverses with respect to addition, therefore the structure can only be called a semiring.

A *tropical monomial* in two variables x and y is a product

$$m = a \odot x^b \odot y^c \quad (a \in \mathbb{R}, b, c \in \mathbb{N})$$

where the powers are computed tropically, too. Considered as a map $m : \mathbb{R}^2 \rightarrow \mathbb{R}$ it represents the linear form $(x, y) \mapsto a \odot x^b \odot y^c = a + bx + cy$.

A *tropical polynomial* is a finite tropical sum of tropical monomials

$$F = a_1 \odot x^{b_1} \odot y^{c_1} \oplus \dots \oplus a_n \odot x^{b_n} \odot y^{c_n}.$$

Again considered as a map $F : \mathbb{R}^2 \rightarrow \mathbb{R}$, F is the piece-wise linear map

$$(x, y) \mapsto \max\{a_1 + b_1x + c_1y, \dots, a_n + b_nx + c_ny\}.$$

As the tropical semiring does not have an inverse operation for addition, it would not make sense to look at zerosets of tropical polynomials. Let us instead recall that the tropical semiring should be thought of as the image of the field K under the map val . Remember that K is defined as the completion of the field $\overline{\mathbb{C}((t))}$ of Puiseux series. As $\text{val}(\overline{\mathbb{C}((t))})$ is dense in \mathbb{R} , we can restrict our considerations to Puiseux series. Let

$$f = p_1(t) \cdot x^{b_1} \cdot y^{c_1} + \dots + p_n(t) \cdot x^{b_n} \cdot y^{c_n} \in \overline{\mathbb{C}((t))}[x, y]$$

be a polynomial and let $(q_1(t), q_2(t)) \in \overline{\mathbb{C}((t))}^2$. Then the sum $f(q_1(t), q_2(t))$ consists of summands corresponding to the monomials $p_i(t) \cdot x^{b_i} \cdot y^{c_i}$ of f , and each summand is a power series in $\overline{\mathbb{C}((t))}$ with a valuation

$$\text{val}(p_i(t) \cdot (q_1(t))^{b_i} \cdot (q_2(t))^{c_i}) = \text{val}(p_i(t)) + b_i \text{val}(q_1(t)) + c_i \text{val}(q_2(t)).$$

Remember that the valuation is given by the negative of the minimal exponent of a power series. A point $(q_1(t), q_2(t)) \in \overline{\mathbb{C}((t))}^2$ can only be in the zeroset $\{f(x, y) = 0\}$, if the monomials sum up to zero, that is, if the minimal exponent of the sum occurs at least twice and these summands cancel. More precisely, a point $(q_1(t), q_2(t))$ can only be in the zeroset $\{f(x, y) = 0\}$, if the maximum

$$\max\{\text{val}(p_1(t)) + b_1 \text{val}(q_1(t)) + c_1 \text{val}(q_2(t)), \dots, \text{val}(p_n(t)) + b_n \text{val}(q_1(t)) + c_n \text{val}(q_2(t))\}$$

is attained by two or more of the terms.

2.6 Definition

Let $f = \sum_i p_i(t) \cdot x^{b_i} \cdot y^{c_i} \in \overline{\mathbb{C}((t))}[x, y]$ be a polynomial. We define its *tropicalization* to be the tropical polynomial

$$\begin{aligned} \text{trop } f &:= \bigoplus_i \text{val}(p_i(t)) \odot x^{b_i} \odot y^{c_i} \\ &= \max_i \{\text{val}(p_1(t)) + b_1 \cdot x + c_1 \cdot y, \dots, \text{val}(p_n(t)) + b_n \cdot x + c_n \cdot y\}. \end{aligned}$$

Note that the condition we checked above to see whether a point $(q_1(t), q_2(t))$ can be in the zeroset of f corresponds to checking whether the maximum $\text{trop } f(\text{val}(q_1(t)), \text{val}(q_2(t)))$ described by the tropical polynomial $\text{trop } f$ evaluated at $(\text{val}(q_1(t)), \text{val}(q_2(t))) \in \mathbb{R}^2$ is attained by two or more monomials.


This observation motivates the following definition:

2.7 Definition

Let F be a tropical polynomial. Then the *tropical curve associated to F* is given by the set of points $(x, y) \in \mathbb{R}^2$ such that the maximum $F(x, y)$ is attained by two or more of the terms (that is the monomials of F). Equivalently, we can say that the tropical curve associated to F is given by the set of points where the piece-wise linear map F is not linear, that is, the tropical curve is the corner locus of the piece-wise linear map F .

2.8 Example


Let $F = 0 \odot x \oplus 0 \odot y \oplus 1 = \max\{x, y, 1\}$ be a tropical polynomial. There are three possibilities how a maximum can be reached by two monomials: either $x = y \geq 1$, or $x = 1 \geq y$, or $y = 1 \geq x$. These three possibilities correspond to three rays, all starting at $(1, 1)$, the first on the diagonal, the second vertically down, the third horizontal. In fact, we get precisely what we suggested to call a tropical line in the beginning: the amoeba of a complex line from far away.




As a second example, let $F = 0 \odot x^2 \oplus 0 \odot y^2 \oplus 0 \odot x \odot y \oplus 1 \odot x \oplus 0 \odot y \oplus 0 = \max\{2x, 2y, x + y, 1 + x, y, 0\}$. Then we have the following possibilities that two terms can both attain the maximum:

- $2x = 2y \geq x + y, 1 + x, y, 0$
 $\Rightarrow x = y \geq 1$ — a diagonal ray starting at $(1, 1)$. But then also $2x = 2y = x + y$ and in fact, three monomials attain the maximum;
- $2x = 1 + x \geq 2y, x + y, y, 0$
 $\Rightarrow x = 1, 1 \geq y$ — a vertical ray starting at $(1, 1)$;
- $2x = y \geq 2y, x + y, 1 + x, 0$
 $\Rightarrow 2x \geq 4x, 3x, 1 + x, 0$ — not possible;
- $2x = 0 \geq 2y, x + y, 1 + x, y$
 $\Rightarrow x = 0, 0 \geq 1 + x = 1$ — not possible;
- $2y = 1 + x \geq 2x, x + y, y, 0$
 $\Rightarrow 2y \geq 4y - 2, 3y - 1, y, 0 \Rightarrow 1 \geq y \geq 0, x = 2y - 1$ — a line segment starting at $(-1, 0)$ and ending at $(1, 1)$;
- $2y = y \geq 2x, x + y, 1 + x, 0$
 $\Rightarrow y = 0, x \leq -1$ — a horizontal ray starting at $(-1, 0)$. But then also $2y = y = 0$ and again three monomials are maximal;
- $x + y = 1 + x \geq 2x, 2y, y, 0$
 $\Rightarrow y = 1, 1 + x \geq 2 \Rightarrow y = 1, x \geq 1, x + 1 \geq 2x \Rightarrow y = x = 1$ — the point $(1, 1)$;
- $x + y = y \geq 2x, 2y, 1 + x, 0$
 $\Rightarrow x = 0, y \geq 2y, 0$ — not possible;
- $x + y = 0 \geq 2x, 2y, 1 + x, y$
 $\Rightarrow y \geq -2y, 2y, 1 - y, y$ — not possible;
- $1 + x = y \geq 2x, 2y, x + y, 0$
 $\Rightarrow y \geq 2y, 0$ — only a point: $x = -1, y = 0$;
- $1 + x = 0 \geq 2x, 2y, x + y, y$
 $\Rightarrow x = -1, 0 \geq y$ — a vertical ray starting at $(-1, 0)$.

That is, the corresponding tropical curve looks like:



Note that unlike the situation in remark 2.1, the primitive integer vectors pointing in the directions of the three edges adjacent to both vertices (the picture should suggest what we mean by using the words “edge” and “vertex”) do not sum up to zero:



However, both the diagonal and the horizontal ray only appear when three monomials attain the maximum, not only two. That is, we might define the weight of an edge, which is one less than the number of monomials taking the maximum along this edge (the general definition will be slightly different, see remark 2.17). Then in our picture in the left, the horizontal ray would be equipped with the weight of 2, and in the right picture, the diagonal ray, too. The weighted sum of the primitive integer vectors will then again be zero. We will see in the following section that this is not a coincidence.

We have now defined tropical curves in two ways: first as the image of a curve over the completion of the field of Puiseux series — a definition which is more helpful to explain why tropical curves should carry most properties of algebraic curves — and second as the corner locus of a tropical polynomial — a definition which allows computations and examples. We also motivated why these two definitions should be equivalent. That they are really equivalent is the result of the following theorem, which is proved for example in [6], Theorem 2.1.1, [25], Theorem 3.3 or [28], Theorem 2.1.

2.9 Theorem (Kapranov’s Theorem)

If $C \subset K^2$ is a curve given by the equation $\{f = 0\}$ and F is the tropical polynomial $F = \text{trop}(f)$, then the tropical curves associated to C (definition 2.4) and associated to F (definition 2.7) coincide.

2.3. A COMBINATORIAL DESCRIPTION OF TROPICAL CURVES

We have now defined tropical curves as images of algebraic curves, and we found an analogous description by means of the tropical semiring. However, both of these descriptions do not give a lot of information about the properties of tropical curves (apart from being piece-wise linear with rational slopes), nor do they describe tropical curves combinatorially. From our examples we expect that tropical curves fulfill for example the condition that the primitive integer vectors of the edges around a vertex sum up to zero. This section will give a combinatorial interpretation of tropical curves that will help to deduce more properties.

2.10 Definition

Let $f = \sum a_i \cdot x^{b_i} \cdot y^{c_i} \in k[x, y]$ be a polynomial (where k now denotes any field, not necessarily the completion K of the Puiseux series from above). The *Newton polygon* $\Delta(f)$ of f is the convex hull of the set $\{(b_i, c_i) | a_i \neq 0\} \subset \mathbb{Z}^2$.


2.11 Definition

Let s be a line segment in \mathbb{R}^2 which starts and ends at an integer valued point. Then the

integer length of s is one less than the number of lattice points on it: the integer length of s is $\#\{s \cap \mathbb{Z}^2\} - 1$.

2.12 Example

Let $f = 2x^2 + y^2 - xy + x + 5$. Then the Newton polygon of f is a triangle with sides of integer length 2:




2.13 Definition


Let $f = \sum a_i \cdot x^{b_i} \cdot y^{c_i} \in k[x, y]$ be a polynomial. Let $D \subset \mathbb{R}^2 \times \mathbb{R}$ be the convex hull of the set $\{(b_i, c_i, a_i) \mid a_i \neq 0\}$. Project the edges of D which can be seen from above (that is, the edges adjacent to faces with an outward pointing normal vector with a positive third coordinate) down to the first factor \mathbb{R}^2 . The image will be a convex subdivision of the Newton polygon, called the *Newton subdivision of f* .

2.14 Example

Let $f = x^2 + 2xy + y^2 + 2x + 2y + 1$. The picture shows the set D .



The projection of the top looks like:



2.15 Theorem


The tropical curve C associated to the tropical polynomial F is dual to the Newton subdivision of F , in the sense that every vertex V of C corresponds to a 2-dimensional polytope of the subdivision and every edge of C is orthogonal to a 1-dimensional polytope.

Furthermore, if a vertex V is adjacent to an edge E , then the 1-dimensional polytope dual to E is in the boundary of the 2-dimensional dual of V .

We will dispense with a precise proof and only give an idea. A proof can for example be found in [23], Proposition 3.11.

2.16 Example

The tropical curve associated to $f = x^2 + 2xy + y^2 + 2x + 2y + 1$ is shown on the left, on the right we show again the Newton subdivision of f determined in example 2.14.




As an idea for the proof of Theorem 2.15, let (b_1, c_1) and (b_2, c_2) be two neighboring points of the Newton subdivision of F . The edge dual to the 1-dimensional polytope which goes from (b_1, c_1) to (b_2, c_2) will be the one that arises when the maximum in $\max\{a_i + b_i x + c_i y\}$ is attained by the two terms $a_1 + b_1 x + c_1 y$ and $a_2 + b_2 x + c_2 y$. When the maximum is attained by these two terms, we have $a_1 + b_1 x + c_1 y = a_2 + b_2 x + c_2 y$, a condition which is fulfilled by the line $y = \frac{b_2 - b_1}{c_1 - c_2} x + \frac{a_2 - a_1}{c_1 - c_2}$. But this line is orthogonal to the line connecting (b_1, c_1) and (b_2, c_2) , which has slope $\frac{c_1 - c_2}{b_1 - b_2}$. The property that (b_1, c_1) and (b_2, c_2) are neighboring helps to see that the condition that the two terms $a_1 + b_1 x + c_1 y$ and $a_2 + b_2 x + c_2 y$ are not only equal but bigger than all other terms is indeed satisfied by some points (x, y) .

2.17 Remark

Note that there can be several tropical polynomials which define the same tropical curve. The way to draw the tropical curve associated to a polynomial F with the aid of the dual Newton subdivision helps us to get an idea why this is true. There can be terms of F such that the corresponding point (b_i, c_i, a_i) in the set D (that we use to determine the Newton subdivision — see 2.13) cannot be seen from above. We can then change the z -coordinate a_i — that is, the corresponding coefficient of the polynomial — up to a maximum without changing the Newton subdivision. In fact, it can also be shown that we can vary this coefficient a_i (up to the maximum) without changing the tropical curve associated to the polynomial. Assume F is a tropical polynomial where we cannot enlarge a coefficient without changing the Newton subdivision (that is, where all points (b_i, c_i, a_i) have their maximal z -coordinate for D). Then the *weight* of an edge e of the tropical curve associated to F is defined to be the number of terms which attain the maximum on the edge e . If F is an arbitrary polynomial, we first have to vary all coefficients up to their maximum before we can compute the weight of an edge of the tropical curve associated to F .


2.18 Remark

Note that duality is not a 1 : 1-correspondence. In fact, many tropical curves can be dual to the same Newton subdivision. The Newton subdivision fixes only the directions in which the edges of the tropical curve point, but not the lengths of the dual edges. This is due to the fact that the coefficients a_i are uniquely determined by the tropical curve, but not by the Newton subdivision corresponding to a polynomial. The picture shows two tropical curves which are dual to the same Newton subdivision.



2.19 Remark

Theorem 2.15 helps us to deduce some nice facts about tropical curves which we assumed to hold already. Let C be a tropical curve, and let V be a vertex of C . Let e_1, \dots, e_n be the edges which start at V . We know that V is dual to a 2-dimensional polytope P , and each edge e_i is dual to a 1-dimensional polytope E_i in the boundary of P . As P is a closed polytope, the vectors E_i sum up to zero. Furthermore, each vector E_i has a certain integer length ω_i . The primitive integral vector u_i pointing in the direction of e_i is orthogonal to E_i . If we multiply u_i with the integer length ω_i , then $\sum_i \omega_i u_i = 0$. That is, we can associate a *weight* to each edge e_i and deduce that the following condition called the *balancing condition* holds at every vertex V of C : the weighted sum of the primitive integer vectors of all edges adjacent to V sum up to 0. (Note that the integer length of a line segment in the Newton subdivision coincides with the number of monomials that attain a maximum together (if all coefficients of the tropical polynomial defining C are at their maximum, see 2.17) — so the definition of weight suggested here coincides with the one suggested in 2.8.) As an example, the following picture shows a tropical curve locally around a vertex:



The primitive integer vectors pointing in direction of the four edges are $\begin{pmatrix} -1 \\ 0 \end{pmatrix}$, $\begin{pmatrix} 0 \\ -1 \end{pmatrix}$, $\begin{pmatrix} 2 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$. The weights are 3, 2, 1 and 1. We can see that the balancing condition holds:

$$3 \cdot \begin{pmatrix} -1 \\ 0 \end{pmatrix} + 2 \cdot \begin{pmatrix} 0 \\ -1 \end{pmatrix} + \begin{pmatrix} 2 \\ 1 \end{pmatrix} + \begin{pmatrix} 1 \\ 1 \end{pmatrix} = 0.$$

The observations we made suggest that we could probably define tropical curves purely combinatorially — roughly as weighted graphs in \mathbb{R}^2 which fulfill the balancing condition. We will indeed define so-called parametrized tropical curves in 4.10 in that way. As this definition carries some properties similar to stable maps, we want to give an introduction to classical enumerative geometry first.

3. CLASSICAL CONCEPTS FOR THE ENUMERATION OF PLANE CURVES

The aim of this chapter is to give an overview about the known results of classical enumerative geometry which we want to revisit tropically later.

3.1 Notation

We fix the complex numbers as ground field for the whole chapter. By \mathbb{P}^2 we denote the complex projective plane, i.e. $\mathbb{P}_{\mathbb{C}}^2$.

Before we start with the definition of the moduli space of stable maps, we would like to give some basic examples about enumerative geometry and point out some problems there are to deal with. As already said in the introduction, enumerative geometry deals with the enumeration of geometric objects — in general, curves — that satisfy certain incidence conditions, like passing through a subspace, or having a given contact order to a subspace. The conditions have to be chosen in such a way that only finitely many of the geometry objects satisfy them.

The easiest example for an enumerative problem is the question: how many lines pass through two given points? The answer is of course 1 — at least if the two points are different. (Here we can see that we have to require the incidence conditions we choose to be in *general position*). Let us prove the answer for the case of projective lines in the plane. A line L in \mathbb{P}^2 is given by a linear polynomial in the three coordinates, $L = \{ax+by+cz = 0\}$. That is, it is again given by three parameters. We can choose another copy of \mathbb{P}^2 , denoted with $\mathbb{P}_{(a:b:c)}^2$, with coordinates $(a : b : c)$, that parametrizes lines in \mathbb{P}^2 . Given a point p_1 in \mathbb{P}^2 , the condition that a line L passes through p_1 is fulfilled by a 1-dimensional linear subspace of the space of lines, that is, by a line in $\mathbb{P}_{(a:b:c)}^2$. The same holds of course for the condition that a line passes through a second point p_2 . If the two points are distinct, then also the two lines in $\mathbb{P}_{(a:b:c)}^2$ are different and meet in one point, which corresponds to the one line that passes through both points.

Let us try to apply the same idea to the next more complicated enumerative question: how many smooth conics in \mathbb{P}^2 pass through five given points?

A conic C in \mathbb{P}^2 is given by a homogeneous polynomial of degree 2:

$$C = \{ax^2 + bxy + cy^2 + dxz + eyz + fz^2 = 0\}.$$

That is, we can again parametrize conics in \mathbb{P}^2 by another projective space, this time by a \mathbb{P}^5 with coordinates $(a : b : c : d : e : f)$. The condition that the conic passes through a point is a linear condition on the coefficients $(a : b : c : d : e : f)$. Therefore we get again a linear subspace of \mathbb{P}^5 of codimension 1, that parametrizes conics passing through that point. The intersection of the five subspaces induced by the five points is one point — again because we chose the points to be in general position, which then also leads to subspaces in general position. The point in the intersection corresponds to the one conic that passes through the five points. The proof seems to be correct so far — but actually, one important argument is missing. We said that a polynomial $ax^2+bxy+cy^2+dxz+eyz+fz^2$ determines a conic — but it does not have to be a smooth conic, it can also be two lines, or even a double line. We have to specify in our reasoning above why we counted smooth conics

and not two lines. This is again true because we chose the points in general position. It is not possible to consider instead of \mathbb{P}^5 the subset which parametrizes only smooth conics, leaving the points which correspond to a union of lines out. This would lead to a non-compact space, but we need compactness for the intersection products.

This example illustrates the problems enumerative geometry has to deal with. A strategy of counting objects is to construct a moduli space which parametrizes these objects. In general, such a moduli space will not be compact. So we have to worry about a good compactification, before we can intersect subvarieties corresponding to objects that fulfill the incidence conditions. At last, we have to check if our result really counted the objects we wanted, or if there is some unwanted contribution from the boundary of the moduli space.

Note that an analogous argument as above does not work for curves of higher degree: it is still true that $\mathbb{P}^{(d+3)d/2}$ parametrizes curves of degree d , but in general, we want to count curves of a fixed genus and degree. The space $\mathbb{P}^{(d+3)d/2}$ is only helpful if we want to count curves of the maximal genus. (Note that lines and conics always have genus 0.)

The moduli space of stable maps is a compact moduli space which seems to be appropriate for many enumerative problems. The basic idea is a different understanding of a curve. We can think of a curve as an embedded object in the surrounding space, or we can think of it as the image of a map from an abstract curve to the space. The moduli space of stable maps is based on the second understanding of curve. Another important feature is the chosen compactification of the space of maps — the stable maps — that turned out to be very helpful for many enumerative questions.

3.1. THE MODULI SPACE OF STABLE MAPS

As already said, we want to understand a curve in a surrounding space as the image of an abstract curve. So before we can deal with maps from abstract curves, we have to study abstract curves themselves. Later on, the images of the abstract curves will be required to meet certain subvarieties. We will therefore provide the abstract curves with certain markings, that shall later map to the given subvarieties.

3.2 Definition

An *n*-marked curve is a tuple (C, x_1, \dots, x_n) where C is a smooth curve and the x_i are distinct points on C . An (iso-)morphism of *n*-marked curves $(C, x_1, \dots, x_n) \rightarrow (C', x'_1, \dots, x'_n)$ is an (iso-)morphism $\varphi : C \rightarrow C'$ satisfying $\varphi(x_i) = x'_i$ for all i .

Let $n \geq 3$ if $g = 0$, and $n \geq 1$ if $g = 1$. $\mathcal{M}_{g,n}$ is the set of isomorphism classes of *n*-marked curves of genus g .

3.3 Example


As an easy example, let us start with rational curves. Of course, there is (up to isomorphism) only one smooth rational curve, namely \mathbb{P}^1 . If we mark a rational curve with 3 markings, there is a unique isomorphism to \mathbb{P}^1 sending the 3 markings to 0, 1 and ∞ . That is, $\mathcal{M}_{0,3}$ is a point, which corresponds to \mathbb{P}^1 with the 3 markings 0, 1 and ∞ . If we mark a rational curve with 4 markings, then we can as before send the first 3 markings to

0, 1 and ∞ . The isomorphism class of this curve is therefore determined by the image of the fourth marking under this isomorphism. This image can be any point in \mathbb{P}^1 , except 0, 1 and ∞ , as we required the marked points to be different. Therefore, $\mathcal{M}_{0,4}$ is equal to $\mathbb{P}^1 \setminus \{0, 1, \infty\}$.

Of course, our argument is on a purely set-theoretic level so far. It can be shown that $\mathbb{P}^1 \setminus \{0, 1, \infty\}$ is a moduli space for isomorphism classes of rational 4-marked curves, but we will only present the ideas here and refer to [10] or [18] for more detailed explanations and proofs.

The question is now how to compactify this space. Let us move towards one of the missing points, for example 0, and check what the corresponding curves look like. The point p in $\mathbb{P}^1 \setminus \{0, 1, \infty\}$ corresponds to a rational curve, a \mathbb{P}^1 , with markings $x_1 = 0$, $x_2 = 1$, $x_3 = \infty$ and $x_4 = p$. When we let p move towards 0 — that is, the corresponding curve has the markings $(0 : 1 : \infty : p)$ — the limit for $p \rightarrow 0$ seems to be a curve with a double marking at 0. But we required the markings to be different. A reason for this is given by the following: the curve with the markings $(0 : 1 : \infty : p)$ is isomorphic to the curve with the markings $(0 : \frac{1}{p} : \infty : 1)$. But if we let p go to 0 here, the limit seems to be a curve with a double marking at ∞ . So the limit is not uniquely defined that way. We can therefore not allow double markings.

Instead, we replace the double marking with a new component carrying the two markings. That is, the curve corresponding to 0 is a reducible curve with two rational components, one carrying the markings x_1 and x_4 , the other carrying the markings x_2 and x_3 . This curve is unique up to isomorphism, as each component has 3 special points, the two markings, and the intersection of the two components. (These 3 special points can as before be mapped uniquely to $0, 1$ and $\infty \in \mathbb{P}^1$, which shows us that the curve is unique up to isomorphism.) Analogously, the two other “missing points” represent reducible curves, but with a different arrangement of the markings.



That is, $\overline{\mathcal{M}}_{0,4}$, the compactification of the space of abstract rational 4-marked curves, is isomorphic to \mathbb{P}^1 . Each $p \in \mathbb{P}^1$ which is not equal to 0, 1 or ∞ represents a smooth rational curve, isomorphic to \mathbb{P}^1 with the markings 0, 1, ∞ and p . 0 represents the reducible curve in the picture on the left, 1 the picture in the middle, and ∞ on the right.

In general, the idea how to compactify the space of smooth abstract n -marked curves of genus g is to allow reducible curves. They have to satisfy the notion of *stability*:

3.4 Definition

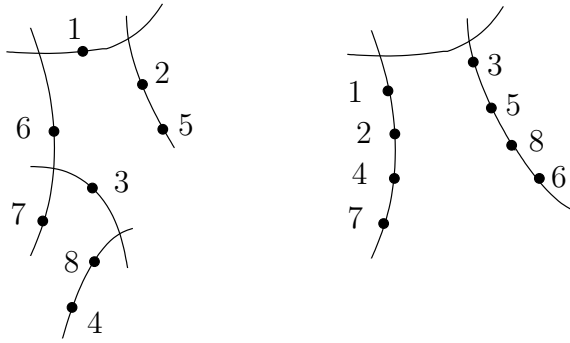
A *pre-stable* n -marked curve is a tuple (C, x_1, \dots, x_n) where C is a connected nodal curve and x_1, \dots, x_n are distinct smooth points on C . An (iso-)morphism of pre-stable n -marked curves $(C, x_1, \dots, x_n) \rightarrow (C', x'_1, \dots, x'_n)$ is an (iso-)morphism $\varphi : C \rightarrow C'$ satisfying $\varphi(x_i) = x'_i$ for all i .

A pre-stable n -marked curve is called *stable*, if its group of automorphisms is finite. This is equivalent to requiring that every rational component contains at least 3 special points (nodes or markings) and every elliptic component contains at least one special point.

Let $n \geq 3$ if $g = 0$, and $n \geq 1$ if $g = 1$. $\overline{\mathcal{M}}_{g,n}$ is defined to be the set of isomorphism classes of stable n -marked curves of genus g and called the *moduli space of stable curves*. It contains $\mathcal{M}_{g,n}$ as a subset, more precisely, this is the subset of smooth curves.

3.5 Example

Here is a picture of two rational 8-marked curves. The one on the left is stable. The one on the right is not stable, because the middle component contains only 2 special points — the 2 nodes.



As rational curves will be of a special interest in section 3.2 and in chapter 7, we cite the following result:

3.6 Theorem

For each $n \geq 3$, $\overline{\mathcal{M}}_{0,n}$ is a projective variety of dimension $n - 3$ and a fine moduli space for stable n -marked rational curves. Especially, $\overline{\mathcal{M}}_{0,4}$ is isomorphic to \mathbb{P}^1 .

As before, see [10] or [18] for more detailed explanations, or [17] for a proof. Here, we just want to give an idea why the dimension is $n - 3$: this is again due to the fact that 3 points on a \mathbb{P}^1 can be sent by an automorphism to 0, 1 and ∞ . Each further marking enlarges the dimension by 1.

3.7 Remark

For higher genus, $\overline{\mathcal{M}}_{g,n}$ should be thought of as a stack, due to the presence of automorphisms. This structure is still good enough to apply intersection theory in some sense. For more information about stacks, see for example [7].

We are now ready to define stable maps.

3.8 Definition

An *pre-stable n -marked map* is a triple (C, x_1, \dots, x_n, f) , where C is a connected nodal curve, x_1, \dots, x_n are distinct smooth points on C and $f : C \rightarrow X$ is a map from C to some ambient space X . An (iso-)morphism of n -marked maps $(C, x_1, \dots, x_n, f) \rightarrow (C', x'_1, \dots, x'_n, f')$ is an (iso-) morphism $\varphi : C \rightarrow C'$ such that $\varphi(x_i) = x'_i$ for all i and such that the following diagram commutes:

$$\begin{array}{ccc}
 C & \xrightarrow{\varphi} & C' \\
 f \searrow & & \swarrow f' \\
 & X &
 \end{array}$$

The class of a pre-stable n -marked map (C, x_1, \dots, x_n, f) is defined to be the element $f_*[C] \in H_2^+(X)$, where $H_2^+(X)$ denotes the space of homology classes of algebraic curves.

A pre-stable n -marked map is called *stable* if its automorphism group is finite.


This is equivalent to requiring that every rational component which is mapped to a point contains at least 3 special points (nodes or marks), and every elliptic component which is mapped to a point contains at least one special point.

For a projective variety X , $g, n \in \mathbb{N}$ and $\beta \in H_2^+(X)$, $\overline{\mathcal{M}}_{g,n}(X, \beta)$ is defined to be the set of isomorphism classes of stable n -marked maps, where C is a curve of genus g , $f : C \rightarrow X$ is a map to X , and the class $f_*[C] \in H_2^+(X)$ is equal to β . It is called the *moduli space of stable maps*.


The subset of stable n -marked maps where the underlying curve C is smooth is called $\mathcal{M}_{g,n}(X, \beta)$.

3.9 Example

The picture shows a 3-marked rational map.



The middle component is contracted to a point, the other two components are mapped to the two components of the image curve shown on the right. The map is stable, because the only component which is contracted to a point has 3 special points — 2 nodes and 1 marking.



This picture shows a 1-marked rational map. The left component is mapped to a point, the right component is mapped to the image shown on the right. This map is not stable, because the component which is contracted to a point contains only 2 special points — 1 node and 1 marking.

3.10 Remark

Note that for a stable map (C, x_1, \dots, x_n, f) the underlying curve C does not need to be a stable curve itself. For example, the automorphism group of a non-constant map $f : \mathbb{P}^1 \rightarrow X$ is finite, whereas the automorphism group of \mathbb{P}^1 is not.

3.11 Theorem

For a smooth projective variety X , $\beta \in H_2^+(X)$ and $g, n \in \mathbb{N}$, $\overline{\mathcal{M}}_{g,n}(X, \beta)$ is a (Deligne-Mumford) stack of expected dimension $-K_X \cdot \beta + (\dim X - 3)(1 - g) + n$.

For a proof, see [3], Theorem 3.14. Here, we only want to present an idea for the dimension count in the case of rational curves mapped to \mathbb{P}^2 : Note first that the homology classes of \mathbb{P}^2 are given by the degrees of the curves. That is, instead of a class β we can fix a degree d . Let $(C, x_1, \dots, x_n, f) \in \overline{\mathcal{M}}_{0,n}(\mathbb{P}^2, d)$. Assume (C, x_1, \dots, x_n, f) is an interior point, that is, C is smooth and therefore isomorphic to \mathbb{P}^1 . That is, f is (up to isomorphism) a map from $\mathbb{P}^1 \rightarrow \mathbb{P}^2$, whose image is a curve of degree d . Such a map is given by 3 homogeneous polynomials in two variables of degree d . The dimension of the space of degree d homogeneous polynomials in two variables is $d+1$. As we need 3 such polynomials, we have $3d + 3$ degrees of freedom to choose them. We have to subtract 1, because two sets of 3 polynomials of degree d give the same map if they differ by a constant factor. Furthermore, we have to subtract 3 for the dimension of the automorphism group of \mathbb{P}^1 . Finally, we have to add n for the markings. Altogether, this dimension count gives us $3d - 1 + n$ which is equal to $-K_{\mathbb{P}^2} \cdot d - 1 + n$, as predicted by theorem 3.11.

The following morphisms of $\overline{\mathcal{M}}_{g,n}(X, \beta)$ will be needed to put incidence conditions on the curves we want to count.

3.12 Definition

The i -th evaluation map

$$\text{ev}_i : \overline{\mathcal{M}}_{g,n}(X, \beta) \rightarrow X : (C, x_1, \dots, x_n, f) \mapsto f(x_i)$$

evaluates the map at the point $x_i \in C$.

Of course, the definition gives us only a set-theoretic assignment, it is not at all clear that ev_i is indeed a morphism. This is shown in [3], Proposition 5.5.

Now we can explain how $\overline{\mathcal{M}}_{g,n}(X, \beta)$ can be used to count curves in X of class β and genus g passing through given subvarieties E_1, \dots, E_n of X . The idea is to count the subset of stable maps in $\overline{\mathcal{M}}_{g,n}(X, \beta)$ such that $f(x_i) \in E_i$ for all i . To do this, we pull back the classes of the subvarieties E_i by ev_i and intersect all these with the (virtual) fundamental class of $\overline{\mathcal{M}}_{g,n}(X, \beta)$. (Here, we need to work with the virtual fundamental class given by the stack and obstruction theory. The virtual fundamental class has similar properties to a fundamental class. For details concerning the construction of the virtual fundamental class of $\overline{\mathcal{M}}_{g,n}(X, \beta)$, see [2] or [1], for example.) We have to choose the E_i in such a way that the dimension of this intersection product is 0, so that it consists of finitely many points. That is, we determine the number

$$\int_{[\overline{\mathcal{M}}_{g,n}(X, \beta)]} (\text{ev}_1^*(E_1) \cap \dots \cap \text{ev}_n^*(E_n))$$

and hope that it is equal to the number of curves of class β and genus g in X that meet the subvarieties E_1, \dots, E_n . This hope will however not always come true. There may for example be some unwanted contributions from the boundary of the moduli stack. Also, we have to be careful that we do not count for a given curve $C' \subset X$ several stable maps $f : C \rightarrow X$ with $f(C) = C'$.

3.13 Notation

Intersection products of the form

$$\int_{[\overline{\mathcal{M}}_{g,n}(X,\beta)]} (\text{ev}_1^*(E_1) \cap \dots \cap \text{ev}_n^*(E_n))$$

are usually called *Gromov-Witten invariants* of the projective variety X .

The arguments above show that there remains some work to be done before we can interpret the Gromov-Witten invariants as enumerative invariants. For an arbitrary variety X this question is not solved yet. Here, we are mainly interested in Gromov-Witten invariants of \mathbb{P}^2 for which it is known that the Gromov-Witten invariants have the enumerative meaning we hoped for.

Let us set up an enumerative problem for \mathbb{P}^2 . The homology classes of \mathbb{P}^2 are given by the degrees of the curves. Instead of a class β we can therefore fix a degree d and require that $f_*[C]$ is a curve of degree d . Bézout's Theorem tells us that a curve will intersect any subvariety E_i of codimension 1, so such a subvariety does not give a condition. (In the dimension count of the intersection product $\bigcap_i \text{ev}_i^*(E_i)$ a subspace E_i of codimension 1 will also not give a contribution bigger 0.) Therefore we leave the codimension 1 subvarieties of \mathbb{P}^2 out for the moment, and require the curves to meet only points. How many points do we need in order to get a 0-dimensional intersection product? The (virtual) dimension of $\overline{\mathcal{M}}_{g,n}(\mathbb{P}^2, d)$ is $3d + g - 1 + n$ by 3.11. The pullback of a point in \mathbb{P}^2 is of codimension 2 in $\overline{\mathcal{M}}_{g,n}(\mathbb{P}^2, d)$. Therefore we need $3d + g - 1$ points to get a 0-dimensional intersection product.

3.14 Definition

Let $N_{\text{cplx}}^{\text{irr}}(d, g)$ be the *number of irreducible degree d and genus g curves in \mathbb{P}^2 passing through $3d + g - 1$ points in general position.*

We have seen in the beginning that $N_{\text{cplx}}^{\text{irr}}(1, 0) = 1$ and $N_{\text{cplx}}^{\text{irr}}(2, 0) = 1$. Up to now we do actually not know that this is well-defined — that is, that the number of curves through $3d + g - 1$ points does not depend on the points we require the curves to meet. This will follow from 3.16 below, which uses the following theorem:

3.15 Theorem

For all $d, g \in \mathbb{N}$, and with $n = 3d + g - 1$,

$$N_{\text{cplx}}^{\text{irr}}(d, g) = \int_{[\overline{\mathcal{M}}_{g,n}(\mathbb{P}^2, d)]} (\text{ev}_1^*(p_1) \cap \dots \cap \text{ev}_n^*(p_n)),$$

where $p_1, \dots, p_n \in \mathbb{P}^2$ denotes a set of points in general position.

For a proof, see for example [18], chapter 3.5. Here, we only want to give an idea — see remark 3.17.

3.16 Remark

The fact that $N_{\text{cplx}}^{\text{irr}}(d, g)$ can be determined as an intersection product of the form

$$\int_{[\overline{\mathcal{M}}_{g,n}(\mathbb{P}^2, d)]} (\text{ev}_1^*(p_1) \cap \dots \cap \text{ev}_n^*(p_n))$$

tells us that the number $N_{\text{cplx}}^{\text{irr}}(d, g)$ does not depend on the choice of the $n = 3d + g - 1$ points, as long as they are in general position. A collection of points in general position will lead to a collection of general substacks $\text{ev}_i^*(p_i)$, and the intersection product of those does not depend on the special position of the substacks, only of their classes. But as points in \mathbb{P}^2 are equivalent, also the pullbacks are equivalent. The independence of $N_{\text{cplx}}^{\text{irr}}(d, g)$ of the position of the points is therefore an easy consequence of the theory of Gromov-Witten invariants.

3.17 Remark

Why do we in fact count curves, when we consider intersection products in $\overline{\mathcal{M}}_{g,n}(\mathbb{P}^r, d)$? Of course, the image $f(C)$ of a general stable map (C, x_1, \dots, x_n, f) is a curve of the correct genus and class, and with $f(x_i) \in E_i$, so it does also meet the given subvarieties. It only remains to check if there are no repetitions. For example, assume $f(C)$ meets E_i in two different points. Then there are two stable maps whose image is $f(C)$ due to the different possibilities of putting the mark x_i on the same curve. However, we only want to count the single curve $f(C)$. Such repetitions are in fact unavoidable if we allow subvarieties E_i of codimension 1. Bézout's Theorem tells us that a curve of degree d will meet a subvariety E_i of degree e and codimension 1 in $d \cdot e$ points. So if we allow codimension-1-subvarieties as incidence conditions, we need to divide our result by this correcting factor in the end. For a subvariety of codimension 2 or more, the situation is different. A general curve will not intersect such a subvariety at all. If we now force it to meet E_i in a point $f(x_i)$, then it will in general intersect E_i only in this point and not in more points.

3.18 Remark

So far, we considered the moduli space of stable maps where we required the underlying curve C to be connected. It is also possible to drop this requirement and to allow disconnected curves, too. This will in fact be done in section 3.3. Intersection products of the form $\int (\text{ev}_1^*(p_1) \cap \dots \cap \text{ev}_n^*(p_n))$ on this space will count the numbers of possibly disconnected curves of given degree and genus through $3d + g - 1$ points in general position. The latter numbers will be denoted by $N_{\text{cplx}}(d, g)$. The main statements of this section hold for these numbers, too.

3.2. KONTSEVICH'S FORMULA TO DETERMINE $N_{\text{cplx}}^{\text{irr}}(d, 0)$

We have seen that we can compute the numbers $N_{\text{cplx}}^{\text{irr}}(d, g)$ as intersection products on the moduli space of stable maps. But we still do not know how to compute these intersection products.

For rational curves, Kontsevich gave in [19] a recursive formula how to determine these numbers. The only initial value needed for the recursion is the number $N_{\text{cplx}}^{\text{irr}}(1, 0) = 1$ of lines through two given points.

This recursive formula can be derived with the help of certain morphisms between the moduli spaces of stable maps and stable curves, the forgetful morphisms. Roughly, a forgetful morphism does what its name claims: it forgets pieces of the data of a stable map. It may forget certain markings, but it may also forget the whole map and send a stable map (C, x_1, \dots, x_n, f) to its underlying curve (C, x_1, \dots, x_n) . (Of course, if we applied this naive description we would end up with non-stable maps or curves in some cases. The morphism does not only forget, but it also stabilizes the result. The precise definition will be given below in 3.24.) The forgetful map π used for Kontsevich's formula forgets all markings but 4 and the map π is a map from $\overline{\mathcal{M}}_{0,n}(\mathbb{P}^2, d)$ to $\overline{\mathcal{M}}_{0,4}$ which sends a stable map to a stabilization of the underlying curve with only 4 markings kept.

A main argument of Kontsevich's formula is the fact that $\overline{\mathcal{M}}_{0,4}$ is isomorphic to \mathbb{P}^1 by 3.6. In \mathbb{P}^1 , all points are linearly equivalent. Therefore we can pullback two different points and get two different, but equivalent divisors in $\overline{\mathcal{M}}_{0,n}(\mathbb{P}^2, d)$. We choose as points in $\overline{\mathcal{M}}_{0,4}$ that we want to pull back the two boundary points



(as in example 3.3). We will intersect the two pullbacks with some more substacks until we have an intersection product of dimension 0, and then compare the two sides of the equation. We will describe this more precisely in the following. We do not give a precise proof but refer to [19] or [5] instead.

3.19 Remark

Note that the reasoning followed here is only valid for rational curves. A crucial argument is that $\overline{\mathcal{M}}_{0,4}$ is isomorphic to \mathbb{P}^1 , so that all its points are linearly equivalent. There is no analogous statement for higher genus. The methods described here have therefore the disadvantage that they may only be applied for rational curves. The algorithm of Caporaso and Harris which will be explained in the following section, 3.3, does not have this disadvantage. However, it involves also the numbers of curves with higher tangency conditions to a line. Therefore we need a more general moduli space in order to derive this algorithm. Hence the algorithm of Caporaso and Harris can be applied more generally, but its proof is more complicated, and the recursion involves more terms. So both algorithms have their advantages and disadvantages, and in any case, they are both worth to study, also later on in the tropical world.

Before we can define forgetful morphisms, we have to make the notion of stabilization precise.


3.20 Definition

Let (C, x_1, \dots, x_n) be a pre-stable n -marked curve. We can associate to it its *stabilization* $s(C, x_1, \dots, x_n)$ by applying the following procedure:

- every rational component which contains only one special point (which has to be a node if the curve is not irreducible, as it is assumed to be connected) is dropped and
- every rational component which contains only two special points (again, if it is not irreducible, at least one has to be a node) is dropped and the two special points are identified.

3.21 Example

The picture shows 3 pre-stable curves (all the components are assumed to be rational) and their stabilizations:




3.22 Definition


Let (C, x_1, \dots, x_n, f) be a pre-stable n -marked map. We can associate to it its *stabilization* $s(C, x_1, \dots, x_n, f)$ by stabilizing each rational component which is contracted to a point by f as a curve following 3.20.

3.23 Example

The picture shows a pre-stable map with four components. We assume that the right component (drawn with a loop) has genus 1 whereas all other components have genus 0. Both the genus 1 and the middle genus 0 component are contracted to a point by f .



The next picture shows its stabilization. The genus 0 component with only two special points which is contracted to a point has been dropped, identifying the two special points. The genus 1 component is not dropped, even so it is contracted to a point by f , because it contains one special point.



3.24 Definition

There are two types of forgetful maps. A *forgetful map*

$$\pi : \overline{\mathcal{M}}_{g,n}(\mathbb{P}^2, d) \rightarrow \overline{\mathcal{M}}_{g,n'}(\mathbb{P}^2, d)$$

for $n' \leq n$ associates to (C, x_1, \dots, x_n, f) the stabilization (see 3.22) of $(C, x_1, \dots, x_{n'}, f)$. (Of course, we can also forget any subset of $\{x_1, \dots, x_n\}$ different from $\{x_{n'+1}, \dots, x_n\}$.)

A *forgetful map*

$$\pi : \overline{\mathcal{M}}_{g,n}(\mathbb{P}^2, d) \rightarrow \overline{\mathcal{M}}_{g,n'}$$

for $n' \leq n$ associates to (C, x_1, \dots, x_n, f) the stabilization (see 3.20) of the curve $(C, x_1, \dots, x_{n'})$. (As before, any other subset of $\{x_1, \dots, x_n\}$ can also be forgotten in addition to the map f .)

Of course, it is not at all obvious that these maps described here on a set-theoretical level are indeed morphisms. This is the result of the following theorem proved in [3]:

3.25 Theorem

The forgetful maps from definition 3.24 are morphisms of (Deligne-Mumford) stacks.

3.26 Remark

The only important forgetful morphism in this section is the map

$$\pi : \overline{\mathcal{M}}_{0,n}(\mathbb{P}^2, d) \rightarrow \overline{\mathcal{M}}_{0,4}$$

sending (C, x_1, \dots, x_n, f) to the stabilization of (C, x_1, \dots, x_4) (we assume $n \geq 4$). When we only talk of the forgetful map π without specifying what the map forgets, we refer to this map.

3.27 Lemma

The two divisors

$$D_{1,2/3,4} = \pi^*(\text{diagram 1}) \quad \text{and} \quad D_{1,3/2,4} = \pi^*(\text{diagram 2})$$


in $\overline{\mathcal{M}}_{0,n}(\mathbb{P}^2, d)$ ($n \geq 4$) (where π denotes the forgetful map from 3.26) are linearly equivalent.

Proof:

This is an easy consequence from 3.6, as the two divisors in $\overline{\mathcal{M}}_{0,4}$ are linearly equivalent. \square

Before we can study these two linearly equivalent divisors in $\overline{\mathcal{M}}_{0,n}(\mathbb{P}^2, d)$, we need to explain something about the structure of special divisors in $\overline{\mathcal{M}}_{0,n}(\mathbb{P}^2, d)$. Let us for simplicity first begin with divisors in $\overline{\mathcal{M}}_{0,n}$. In example 3.3 we learned that the boundary

of $\overline{\mathcal{M}}_{0,4}$ consists of the 3 points which correspond to reducible curves. We have also seen that we needed reducible curves in order to replace curves where two markings coincide. That is, when two markings move close to each other in a family of curves, the limit is not the same curve with a double marking, but a curve with a new component with the two markings that came together. This construction does not only hold for $\overline{\mathcal{M}}_{0,4}$, but in general. For example, take a family of curves in $\overline{\mathcal{M}}_{0,6}$, where the markings 2 and 5 come together. The family is shown on the left, the limit on the right:



Especially, when we take the closure of the subset of curves with two components, one with markings 1, 4 and 3 and the other with markings 6, 2 and 5, then the curve on the right is contained in this closure, too, as it is the limit of a family contained in this subset.

Now consider the subset of curves with one node. Each such curve has two components, and at least 2 markings on each component. Each further marking can move on the component in a 1-dimensional family and enlarges the dimension of this set by 1. That is, the dimension of the subset of curves with one node is $n - 4$. As $\dim \overline{\mathcal{M}}_{0,n} = n - 3$ by theorem 3.6, the subset of curves with one node is of codimension 1. Analogously, we can see that the subset of curves with δ nodes is of codimension δ .

That is, we can for example take the subset of curves with one node in $\overline{\mathcal{M}}_{0,6}$, and with markings 1, 4 and 3 on one of the components and 6, 2 and 5 on the other component. Its closure is a subset of codimension 1, hence a divisor. Its closure contains the curve in the picture above on the left with 2 nodes, too.

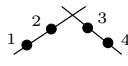
In general if we take the closure of the subset in $\overline{\mathcal{M}}_{0,n}$ of curves with one node and possibly with some requirements on where the markings should lie, this is a divisor.

It also contains curves with more than one node in the boundary.


Also for $\overline{\mathcal{M}}_{0,n}(\mathbb{P}^2, d)$: the subset of stable maps, whose underlying curve has one node and possibly some requirements on where the markings lie, is a divisor. It contains maps whose underlying curve has more than one node in its boundary.

For more detailed explanation about these divisors, see for example [18], chapter 1.5 and 2.7.

Let us now study the two linear equivalent divisors $D_{1,2/3,4}$ and $D_{1,3/2,4}$ in $\overline{\mathcal{M}}_{0,n}(\mathbb{P}^2, d)$ from lemma 3.27. As before, we only give a set-theoretic argument.

Let us determine what stable maps (C, x_1, \dots, x_n, f) in $\overline{\mathcal{M}}_{0,n}(\mathbb{P}^2, d)$ map to  under π . If we forget all markings but x_1, \dots, x_4 and the map, then we end up with a reducible curve with two components, x_1 and x_2 on one component, x_3 and x_4 on the other. That is, also C must have at least 2 components, and x_1 and x_2 must be on one component (possibly after dropping some other unstable components), x_3 and x_4 on the other. That

is, $D_{1,2/3,4}$ is (as a set) equal to the closure of the subset of maps in $\overline{\mathcal{M}}_{0,n}(\mathbb{P}^2, d)$, where the underlying curve has one node, x_1 and x_2 on one component, and x_3 and x_4 on the other. When we take an interior point of this divisor (not a point in the boundary), then the underlying curve has exactly one node and not more. Hence it has two components. As the image $f(C)$ is of degree d , there are several possibilities for the degrees of the images of the two components. The first component can be mapped to a degree d_A curve, but then the image of the second component must be of degree $d_B = d - d_A$. Also, the other markings $\{x_5, \dots, x_n\}$ can lie anywhere on the two components. That is, the divisor $D_{1,2/3,4}$ is reducible. Its components are parametrized by partitions d_A and d_B of d (that is, $d_A + d_B = d$) and disjoint subsets A and B whose union is $\{x_1, \dots, x_n\}$, such that $x_1, x_2 \in A$ and $x_3, x_4 \in B$.



3.28 Notation

For d_A and d_B with $d_A + d_B = d$ and A, B with $A \cup B = \{1, \dots, n\}$ we denote by $D(d_A, d_B, A, B)$ the closure of the subset of stable maps (C, x_1, \dots, x_n, f) in $\overline{\mathcal{M}}_{0,n}(\mathbb{P}^2, d)$ such that

- C has precisely one node,
- all x_i where $i \in A$ are on one component,
- this component is mapped to a curve of degree d_A in \mathbb{P}^2 under f ,
- all x_i where $i \in B$ are on the other component and
- this other component is mapped to a curve of degree d_B in \mathbb{P}^2 under f .

So far we know the support of the divisor $D_{1,2/3,4}$, we do not know if any of the reducible components occur with a higher order on the pullback. But each order is 1, which is the result of the following theorem proved in [18]:

3.29 Theorem

With the notation from 3.28 we have the following equality for the divisor $D_{1,2/3,4}$ from lemma 3.27:

$$D_{1,2/3,4} = \sum_{d_A+d_B=d} \sum_{A,B} D(d_A, d_B, A, B),$$

where the second sum goes over all A and B with $A \cup B = \{1, \dots, n\}$, $1, 2 \in A$ and $3, 4 \in B$. An analogous equality holds for $D_{1,3/2,4}$.

We are now ready for (at least an idea of) Kontsevich's formula.

3.30 Definition

Let $N(d)$ be a collection of numbers given for each $d \in \mathbb{N}$. We say that the numbers $N(d)$

satisfy Kontsevich's formula, if for all $d > 1$ the following recursive formula holds:

$$N(d) = \sum_{\substack{d_A+d_B=d \\ d_A, d_B > 0}} \left(d_A^2 d_B^2 \binom{3d-4}{3d_A-2} - d_A^3 d_B \binom{3d-4}{3d_A-1} \right) N(d_A) N(d_B).$$

3.31 Theorem (Kontsevich's formula)

The numbers $N_{\text{cplx}}^{\text{irr}}(d, 0)$ from definition 3.14 satisfy Kontsevich's formula.

For a proof, see [19], [5] or [18]. Here, we want to present an idea of the proof. As already mentioned, we intersect both sides of the equation $[D_{1,2/3,4}] = [D_{1,3/2,4}]$ in $\overline{\mathcal{M}}_{0,3d}(\mathbb{P}^2, d)$ with $\text{ev}_1^*(L_1) \cap \text{ev}_2^*(L_2) \cap \text{ev}_3^*(p_3) \cap \dots \cap \text{ev}_{3d}^*(p_{3d})$, where the set of lines L_i and points p_i is in general position. That is, we require the first two markings each to meet a line, L_1 (respectively L_2) and all other points to meet a point.

Let us first determine the left side of this equation. From theorem 3.29 we know that

$$\begin{aligned} & [D_{1,2/3,4}] \cap \text{ev}_1^*(L_1) \cap \text{ev}_2^*(L_2) \cap \text{ev}_3^*(p_3) \cap \dots \cap \text{ev}_{3d}^*(p_{3d}) = \\ & \sum_{d_A+d_B=d} \sum_{A,B} D(d_A, d_B, A, B) \cap \text{ev}_1^*(L_1) \cap \text{ev}_2^*(L_2) \cap \text{ev}_3^*(p_3) \cap \dots \cap \text{ev}_{3d}^*(p_{3d}) \end{aligned}$$

where the second sum goes over all A and B with $A \cup B = \{1, \dots, n\}$, $1, 2 \in A$ and $3, 4 \in B$. Let us specify what summands there are: It can for example happen that $d_A = 0$ and $d_B = d$. That is, for all stable maps belonging to this summand, the component which contains the markings x_1 and x_2 is mapped to a point. As x_1 is required to meet L_1 and x_2 to meet L_2 , this point must be $L_1 \cap L_2$. There cannot be another marking on this component, as $p_i \neq L_1 \cap L_2$ for all i . That is, all other markings lie on the second component which is mapped with degree d to \mathbb{P}^2 . The image curve is therefore a curve of degree d passing through the $3d - 1$ points $L_1 \cap L_2, p_3, \dots, p_{3d}$. Hence, the summand

$$\sum_{A,B} D(0, d, \{1, 2\}, \{3, \dots, 3d\}) \cap \text{ev}_1^*(L_1) \cap \text{ev}_2^*(L_2) \cap \text{ev}_3^*(p_3) \cap \dots \cap \text{ev}_{3d}^*(p_{3d})$$

is equal to $N_{\text{cplx}}^{\text{irr}}(d, 0)$.

Next, let us look at a summand with $d_B = 0$. There is actually no stable map in this intersection, as the markings x_3 and x_4 lie on the component which is then contracted to a point. x_3 and x_4 are required to meet distinct points p_3 and p_4 , which is not possible if the component on which they lie is mapped to one single point.

So there are only summands left with $d_A, d_B > 0$. A stable map in such a summand has two components called C_A and C_B , one is mapped to a degree d_A curve, the other to a degree d_B curve. We can consider these two components as distinct stable maps (C_A, x_i, f_{C_A}) where i runs over all $i \in A$, and (C_B, x_i, f_{C_B}) where i runs over all $i \in B$. They are then elements of $\overline{\mathcal{M}}_{0, \#A}(\mathbb{P}^2, d_A)$, respectively $\overline{\mathcal{M}}_{0, \#B}(\mathbb{P}^2, d_B)$. The dimension of these two spaces (see theorem 3.11) tells us that $\#A$ must be equal to $3d_A + 1$ (because the two markings x_1 and x_2 which are required to meet a line are in A , and they do not change the dimension) and $\#B$ to $3d_B - 1$, as otherwise we would not have a 0-dimensional intersection product in total. There are $\binom{3d-4}{3d_A-1}$ possibilities how the markings x_5, \dots, x_{3d} can be arranged on the two components. Then (C_A, x_i, f_{C_A}) is a stable map which meets the two lines L_1 and L_2 . As in remark 3.17 we know that L_1 (respectively, L_2) intersects

$f(C_A)$ in d_A points. Therefore, there are d_A^2 different stable maps which have the same image $f(C_A)$. The image curve $f(C_A)$ is a curve of degree d_A passing through $3d_A - 1$ points. The image curve $f(C_B)$ is a curve of degree d_B passing through $3d_B - 1$ points. The two curves $f(C_A)$ and $f(C_B)$ intersect due to Bézout's theorem in $d_A \cdot d_B$ points. That is, there are $d_A \cdot d_B$ different stable maps with the image $f(C_A) \cup f(C_B)$, as each of the intersection points can be the node where the two components C_A and C_B are attached to each other. That is, a summand with $d_A, d_B > 0$ gives a contribution of $\binom{3d-4}{3d_A-1} \cdot d_A^3 \cdot d_B \cdot N_{\text{cplx}}^{\text{irr}}(d_A, 0) \cdot N_{\text{cplx}}^{\text{irr}}(d_B, 0)$.

Now, we can determine the right side of the equation.

Let us first check the summands where d_A or d_B is 0. Then one component is contracted to a point. But this component contains at least two markings, x_1 and x_3 , respectively x_2 and x_4 . One is required to meet a line, the other to meet a point. But as none of the points lies on the line, this is not possible, if the component is mapped to one point. So there is no contribution from summands where one of the degrees is 0.

So on the right side we only have summands with $d_A, d_B > 0$. Again, a stable map in such a summand has two components C_A and C_B . Considered as two distinct stable maps, they lie in $\overline{\mathcal{M}}_{0, \#A}(\mathbb{P}^2, d_A)$, respectively $\overline{\mathcal{M}}_{0, \#B}(\mathbb{P}^2, d_B)$. One marking on each of the maps is required to meet a line, therefore $\#A = 3d_A$ and $\#B = 3d_B$. There are $\binom{3d-4}{3d_A-2}$ possibilities how the markings x_5, \dots, x_{3d} can be arranged on the two components. The line L_1 intersects $f(C_A)$ in d_A points. So we have again d_A different stable maps mapping to the same image curve. Also, for $f(C_B)$, we have d_B different stable maps mapping to it due to the line L_2 . And as above, there are $d_A \cdot d_B$ possibilities for the node where the two components are attached to each other. That is, altogether a summand on the right side of the equation with $d_A, d_B > 0$ contributes $\binom{3d-4}{3d_A-2} \cdot d_A^2 \cdot d_B^2 \cdot N_{\text{cplx}}^{\text{irr}}(d_A) \cdot N_{\text{cplx}}^{\text{irr}}(d_B)$.

Finally, let us bring the summands with $d_A, d_B > 0$ on one side of the equation. The result is the desired formula.

3.32 Remark

Note that a collection of numbers satisfying Kontsevich's formula is uniquely determined by the initial value $N(1)$. Therefore, theorem 3.31 tells us that we can determine the numbers $N_{\text{cplx}}^{\text{irr}}(d, 0)$ recursively, only using the initial information that there is one line through two points.

3.3. THE ALGORITHM OF CAPORASO AND HARRIS TO DETERMINE $N_{\text{cplx}}(d, g)$

In this section, we want to present another recursive formula that determines the numbers $N_{\text{cplx}}(d, g)$ from remark 3.18 from the only initial value that $N_{\text{cplx}}(1, 0) = 1$. As before, we do not give proofs, but only present some ideas why such a formula should be true. The algorithm is more complicated, because it involves also the numbers of curves that do not only pass through certain given points, but satisfy in addition tangency conditions (of higher order) to a line (see remark 3.19). We need to define a substack of the moduli space of stable maps, where the markings are not only required to meet subspaces, but also to be tangent (of higher order) to a subspace.

The main idea of this algorithm is to change the position of the points that the curves are required to meet. Assume we want to determine the number $N_{\text{cplx}}(d, g)$. We know by theorem 3.11 that we need $3d + g - 1$ points in general position in order to get a finite number of curves passing through them. We fix a line L . Then we specialize the position of the points by moving one after the other to the line L .

An important argument for the algorithm is that the number of curves through the set of points stays the same, even if we specialize the position (as long as it is finite). However, we have to count certain degenerations of the curves (that is, elements in the boundary of our moduli space) that may arise after specializing the points. For example, when we count curves of degree 3 and we specialize the fourth point to lie on L , then there is no irreducible curve of degree 3 through this set of points. The line L splits off as a component, leaving a curve of degree 2 as second component.


These remaining components that arise in the degenerations of our curves may now have a higher tangency order to the line L in some points.

Another important argument of the proof is to determine with which multiplicity the degenerations in the boundary of the moduli space have to be counted, as components of an element in the Chow group of the space of stable maps. In fact, some of these components that parametrize the degenerations have to be counted with a higher multiplicity. More precisely, whenever we have a component of our curve with a point of tangency order m to the line L , we have to multiply with m .


Let us start with an example of the algorithm:

3.33 Example


Assume we want to count rational cubics through 8 points p_1, \dots, p_8 in general position, that is, we want to determine $N_{\text{cplx}}(3, 0)$.



We move the point p_1 to our chosen line L . Next, we move the point p_2 to L . So far, nothing happens, there are no degenerations, but only smooth cubics through these 8 points. In fact, the two points p_1 and p_2 just determine the line L .




Now, let us move the point p_3 to L . There are still smooth cubics that pass through these points. But now, there is also a degeneration that we have to count: it is possible that the line L splits off as a component. The remaining component is then a curve of degree 2 through 5 points. We can assume that this number is recursively known. (In fact, we know from the beginning of the chapter that it is 1.) The smooth conic through the five remaining points intersects the line L — and hence, the other component of the stable map — in two points. Therefore, there are 2 distinct stable maps that have the same image curve. We have to multiply by 2.



Next, we move the point p_4 to L . Now there is no smooth cubic through this set of points, as no smooth cubic can intersect a line in more than 3 points. So the line has to split off as a component, and we have to determine how the remaining component may look like. It is in any case a curve of degree 2 through 4 points. There are several possibilities. (It is part of Caporaso's and Harris' proof to determine which possibilities actually arise.)

- It could be a conic tangent to L .


This is now where the higher tangency orders can appear: we know that a smooth conic is determined by 5 given points. But alternatively, it can also be determined by 4 points and a line to which it is tangent. (This follows by a dimension count: the condition to pass through a point is of the same dimension than that to be tangent to a given line.) We can again assume that this number is given recursively, in fact, it is equal to 2. As mentioned in the beginning, we have to multiply this number by 2 due to the component with tangency order 2 to the line.



$$2 \cdot 2$$


- It could be a conic that passes through one of the other 3 points that was already moved to the line L .

The number of conics through 5 points is known to be 1. Contrary to the step before, where we also had a conic through 5 points, there is now no choice for the node of the stable map mapping to this curve, as the markings are required to be different from the nodes. That is, the node is fixed to be on the intersection point of the conic and L different from the 3 points p_1 , p_2 and p_3 on L . We have to multiply the number of conics through 5 points with 3 due to the 3 possibilities for the point on L through which the conic passes.



- It could be a reducible curve, that is, a union of two lines.

Of course, the degree 2 curve that splits off does not need to be smooth. It can also be a union of two lines which are determined by the four points. The number of lines through 2 points is known to be 1. But we have to multiply with the factor 3, as there are 3 possibilities to distribute the four points on the two lines.



Let us count all the degenerations whose numbers are assumed to be known recursively: we have 2 for the conic through 5 points and the line. In the next step, there are 2 for the conic through 4 points and tangent to a line — but we also have to multiply with another factor of 2 here, since we are counting curves with a higher tangency order to the line, so altogether 4. Furthermore 3 for the conics through the 4 points and through one more point on L , and finally 3 possibilities for the two lines, each counting once, as there is one line through two points. Altogether, we have $2 + 4 + 3 + 3 = 12$ rational cubics through 8 points.

Let us now give a more formal description of the algorithm of Caporaso and Harris. Before we can do so, we have to fix some notational conventions:

3.34 Notation

Let $\alpha = (\alpha_1, \alpha_2, \dots)$ be a *finite sequence* of natural numbers, that is, almost all α_k are zero. If $\alpha_k = 0$ for all $k > n$ we will also write this sequence as $\alpha = (\alpha_1, \dots, \alpha_n)$. For two sequences α and β we define

$$\begin{aligned}
 |\alpha| &:= \alpha_1 + \alpha_2 + \dots, \\
 I\alpha &:= 1\alpha_1 + 2\alpha_2 + 3\alpha_3 + \dots, \\
 I^\alpha &:= 1^{\alpha_1} \cdot 2^{\alpha_2} \cdot 3^{\alpha_3} \cdot \dots, \\
 \alpha + \beta &:= (\alpha_1 + \beta_1, \alpha_2 + \beta_2, \dots), \\
 \alpha \geq \beta &: \Leftrightarrow \alpha_n \geq \beta_n \text{ for all } n, \text{ and}
 \end{aligned}$$

$$\binom{\alpha}{\beta} := \binom{\alpha_1}{\beta_1} \cdot \binom{\alpha_2}{\beta_2} \cdot \dots$$

We denote by e_k the sequence which has a 1 at the k -th place and zeros everywhere else.

In example 3.33 we have seen that we do not only need the numbers $N_{\text{cplx}}(d, g)$ in the recursion, but also the numbers of curves satisfying in addition tangency conditions to the line L . The following definition makes this notion precise:

3.35 Definition

Let $d \geq 0$ and g be integers, and let α and β be two sequences with $I\alpha + I\beta = d$. Fix a line $L \subset \mathbb{P}^2$. Then we denote by $N_{\text{cplx}}^{\alpha, \beta}(d, g)$ the number of nodal, not necessarily irreducible curves of degree d and genus g that

- intersect L in α_i fixed general points of L with contact order i for all $i \geq 1$,
- intersect L in β_i arbitrary points of L with contact order i for all $i \geq 1$ and
- pass additionally through $2d + g + |\beta| - 1$ more points in general position in \mathbb{P}^2 .

That is, $N_{\text{cplx}}^{\alpha, \beta}(d, g)$ is the number of degree d and genus g curves that have the given contact orders to L (in some fixed, some arbitrary points). Note that this definition generalizes the definition of 3.18 from before: $N_{\text{cplx}}(d, g) = N_{\text{cplx}}^{(0), (d)}(d, g)$ in the new notation.

3.36 Remark

Once we defined an appropriate moduli stack parametrizing curves with the properties from definition 3.35, we will see for dimensional reasons why the number $2d + g + |\beta| - 1$ of points that the curves are required to meet is the correct number to choose in order to get a finite number of curves (see definition 3.42 and theorem 3.43).

The result of Caporaso and Harris is now that these numbers (and with them, also the numbers $N_{\text{cplx}}(d, g)$) can be computed recursively:

3.37 Definition

Let $N^{\alpha, \beta}(d, g)$ be a collection of numbers given for each $d \in \mathbb{N}$, $g \in \mathbb{Z}$ and finite sequences α and β with $I\alpha + I\beta = d$. We say that this collection satisfies the Caporaso-Harris formula if

$$\begin{aligned} N^{\alpha, \beta}(d, g) = & \sum_{k | \beta_k > 0} k \cdot N^{\alpha + e_k, \beta - e_k}(d, g) \\ & + \sum I^{\beta' - \beta} \cdot \binom{\alpha}{\alpha'} \cdot \binom{\beta'}{\beta} \cdot N^{\alpha', \beta'}(d - 1, g') \end{aligned}$$

for all d, g, α and β as above with $d > 1$, where the second sum is taken over all α', β' and g' satisfying

$$\begin{aligned} \alpha' &\leq \alpha, \\ \beta' &\geq \beta, \\ I\alpha' + I\beta' &= d - 1, \\ g - g' &= |\beta' - \beta| - 1, \text{ and} \\ d - 2 &\geq g - g'. \end{aligned}$$

3.38 Theorem (The algorithm of Caporaso and Harris)

The numbers $N_{\text{cplx}}^{\alpha,\beta}(d, g)$ of definition 3.35 satisfy the Caporaso-Harris formula.

3.39 Remark

The result 3.38 is originally shown in [4]. It is proved there with different methods. The moduli space of stable maps is not used there, but Hilbert schemes. Both languages have their advantages and disadvantages (see [32], remark 1.1). In the following, we want to present some ideas of the proof, sticking to the language of stable maps in order to stay consistent with the previous section. The results of [4] were reproved (and generalized) in the language of stable maps in [32], for example. Also, a more general result (although restricted to rational curves) is shown in [11] by means of stable maps.

3.40 Remark

Note that a collection of numbers satisfying the Caporaso-Harris formula is uniquely determined by the initial values for $d = 1$. In particular, theorem 3.38 tells us that we can compute the numbers $N_{\text{cplx}}^{\alpha,\beta}(d, g)$ recursively, starting with the initial value that there is one line through two given points.

Let us now give some ideas for the proof of theorem 3.38. First, we have to take a slightly different moduli space:

3.41 Definition

A *quasi-stable map* has the same properties as a stable map, the only difference is, that the underlying curve is allowed to be disconnected. All constructions that work for stable maps (the moduli stack, divisors, evaluation morphisms, forgetful morphisms and so on) work for quasi stable maps in the same way. The *moduli stack of quasi-stable maps* is denoted with $\overline{\mathcal{M}}'_{g,n}(\mathbb{P}^2, d)$.

3.42 Definition

Fix a line L in \mathbb{P}^2 . Let $d \geq 0$ and g be given, and furthermore sequences $\alpha = (\alpha_1, \dots, \alpha_m)$ and $\beta = (\beta_1, \dots, \beta_{m'})$ satisfying $I\alpha + I\beta = d$. Let $\Gamma = \{q_{1,1}, \dots, q_{1,\alpha_1}, \dots, q_{m,1}, \dots, q_{m,\alpha_m}\} \subset L$ be a set of general points on L .

Then we denote by $V_n^{\Gamma,(\alpha,\beta)}(d, g)$ the closure in $\overline{\mathcal{M}}'_{g,|\alpha|+n}(\mathbb{P}^2, d)$ of the subset of quasi stable maps

$$(C, x_{1,1}, \dots, x_{1,\alpha_1}, \dots, x_{m,1}, \dots, x_{m,\alpha_m}, x_1, \dots, x_n, f)$$

(note that we choose this at first glance complicated labelling of the markings, because the markings labelled x_{i,j_i} will be required to be points of contact order i to L , whereas the markings labelled x_i will as usually be required to meet points)

satisfying

- each connected component of C maps birationally onto its image in \mathbb{P}^2 ,
- $f^{-1}(L)$ is a finite set

$$f^{-1}(L) = \{x_{1,1}, \dots, x_{1,\alpha_1}, \dots, x_{m,1}, \dots, x_{m,\alpha_m}\} \cup \{r_{1,1}, \dots, r_{1,\beta_1}, \dots, r_{m',1}, \dots, r_{m',\beta_{m'}}\}$$

of $|\alpha| + |\beta|$ smooth points of C ,

- the divisor $f^*(L)$ on C is given by

$$f^*(L) = \sum_{i=1, \dots, m} \sum_{j=1, \dots, \alpha_i} i \cdot x_{i,j_i} + \sum_{i=1, \dots, m'} \sum_{j=1, \dots, \beta_i} i \cdot r_{i,j_i}$$

and

- $f(x_{i,j_i}) = q_{i,j_i}$ for all $i = 1, \dots, m$ and $j_i = 1, \dots, \alpha_i$.

This is a map-theoretic equivalent of the Severi variety in [4].

3.43 Theorem

$V_n^{\Gamma, (\alpha, \beta)}(d, g)$ is a substack of dimension $2d + g + |\beta| - 1 + n$.

For a proof, see [32], Theorem 3.1.

3.44 Remark

Similar to the proof of Kontsevich's formula, the idea of the proof of theorem 3.38 is to determine the number of points in the intersection product

$$\int_{[V_n^{\Gamma, (\alpha, \beta)}(d, g)]} (\text{ev}_1^*(p_1) \cap \dots \cap \text{ev}_n^*(p_n)),$$

where p_1, \dots, p_n is a set of points in general position in \mathbb{P}^2 . The number n is chosen to be equal to $2d + g + |\beta| - 1$, in order to get indeed a 0-dimensional intersection. Analogously to what we have seen in section 3.2 (see for example remark 3.17), it is not at all obvious that the quasi stable maps we count then are actually in bijection to the image curves through the appropriate number of points. (Of course, the image curve will be a curve that fulfills the requirements of definition 3.35). That is, also if we are able to determine this intersection product, we still have to worry about the enumerative interpretation of this result. We refer to [32], chapter 4 for a discussion of this question. The result presented there is that for \mathbb{P}^2 (actually, it is shown for more general surfaces than \mathbb{P}^2) the numbers we get are enumerative, that is, by computing this intersection product, we actually count the numbers $N_{\text{cplx}}^{\alpha, \beta}(d, g)$ of definition 3.35.

3.45 Notation

Intersection products of the form

$$\int_{[V_n^{\Gamma, (\alpha, \beta)}(d, g)]} (\text{ev}_1^*(p_1) \cap \dots \cap \text{ev}_n^*(p_n)),$$

or more generally, intersection products that should count curves satisfying conditions of higher contact order to a given subspace in addition to meeting certain given points, are called *relative Gromov-Witten invariants*.

So let us next try to compute this intersection product.

As in remark 3.16, the divisor class of $\text{ev}_i^*(p_i)$ is independent of p_i .

Therefore we can choose a point $q \in L$ not in Γ and intersect $V_n^{\Gamma, (\alpha, \beta)}(d, g)$ with $\text{ev}_1^*(q)$. We know that we can make this choice for q , as the class of $\text{ev}_1^*(q)$ does not depend on q . We want to describe the components of this intersection again in terms of other substacks $V_{n'}^{\Gamma', (\alpha', \beta')}(d', g')$. Once we understand the components of $V_n^{\Gamma, (\alpha, \beta)}(d, g) \cap \text{ev}_1^*(q)$ in terms

of some other $V_{n'}^{\Gamma',(\alpha',\beta')}(d',g')$, we can intersect this substack with the remaining point conditions $(\text{ev}_2^*(p_2) \cap \dots \cap \text{ev}_n^*(p_n))$ to get a 0-dimensional set that we can count.

The equality

$$\begin{aligned} N_{\text{cplx}}^{\alpha,\beta}(d,g) &= V_n^{\Gamma,(\alpha,\beta)}(d,g) \cap (\text{ev}_1^*(p_1) \cap \dots \cap \text{ev}_n^*(p_n)) \\ &= (V_n^{\Gamma,(\alpha,\beta)}(d,g) \cap \text{ev}_1^*(q)) \cap (\text{ev}_2^*(p_2) \cap \dots \cap \text{ev}_n^*(p_n)) \end{aligned}$$

leads to the recursive relations we want, when we insert the result of our study of the intersection $(V_n^{\Gamma,(\alpha,\beta)}(d,g) \cap \text{ev}_1^*(q))$.

There are two types of components in this intersection.

First, there is a set K_1 of components consisting of maps where one of the points r_{i,j_i} is mapped to q . That is, $x_1 = r_{i,j_i}$ and is now one of the points that is mapped to L . (In example 3.33, in the step when we moved the third point p_3 to the line L — which corresponds to intersecting with $\text{ev}_1^*(p_3)$ for the point $p_3 \in L$ — the smooth cubic through the 3 points on L is for instance an element of such a component of first type.)

Second, there is a set K_2 of components consisting of maps where one component is mapped onto the line L . Let $(C, \{x_{i,j_i}\}, x_1, \dots, x_n, f)$ be a general point of one component of K_2 .

That is, $C = C' \cup C''$, $x_1 \in C'$ and $f|_{C'}(C') = L$. It can be shown that indeed $f|_{C'_*}(C') = 1 \cdot L$ (see [32], proof of Theorem 5.1). $f|_{C''}(C'')$ is then a curve of degree $d - 1$.

There are several possibilities how the markings can be arranged on C' respectively C'' . Any subset of the $\{x_{i,1}, \dots, x_{i,\alpha_i}\}$ can lie on C'' . For all i , let us assume the subset of $\{x_{i,1}, \dots, x_{i,\alpha_i}\}$ of points on C'' consists of α'_i elements with $\alpha'_i \leq \alpha_i$. Altogether, the subset of $\{x_{1,1}, \dots, x_{1,\alpha_1}, \dots, x_{m,1}, \dots, x_{m,\alpha_m}\}$ of points on C'' then has $|\alpha'|$ elements. Let us denote this subset by $\{x'_{1,1}, \dots, x'_{1,\alpha'_1}, \dots, x'_{m,1}, \dots, x'_{m,\alpha'_m}\}$. There are $\binom{\alpha}{\alpha'}$ possibilities for the choice of such a subset. The subset induces a subset $\Gamma' \subset \Gamma$ of the images of these markings.

Assume the markings $x_2, \dots, x_{n'}$, $n' \leq n$, lie on C'' .

Our aim is now to relate the map

$$(C'', x'_{1,1}, \dots, x'_{1,\alpha'_1}, \dots, x'_{m,1}, \dots, x'_{m,\alpha'_m}, x_2, \dots, x_{n'}, f|_{C''})$$

to an element of a substack of the form $V_{n'}^{\Gamma',(\alpha',\beta')}(d-1,g')$ in $\overline{\mathcal{M}}'_{g,|\alpha'|+n'}(\mathbb{P}^2, d-1)$. For this, we study the pullback of the line L on C'' as a divisor.

Of course, the pullback contains the markings $\{x'_{1,1}, \dots, x'_{1,\alpha'_1}, \dots, x'_{m,1}, \dots, x'_{m,\alpha'_m}\}$, the multiplicity of $(f|_{C''})^*(L)$ at a point x'_{i,j_i} is i and $f|_{C''}(x'_{i,j_i})$ is the corresponding point in Γ' .

The set of points in $C' \cap C''$ is mapped to the line L as C' is. Therefore, also these points are contained in the divisor $(f|_{C''})^*(L)$. Let γ be the sequence defined by the multiplicity of $(f|_{C''})^*(L)$ at these points. That is, if $\gamma = (\gamma_1, \gamma_2, \dots, \gamma_l)$ there are γ_1 points of multiplicity 1, γ_2 of multiplicity 2, and so on in $C' \cap C''$.

By taking a family of maps around $(C, \{x_{i,j_i}\}, x_1, \dots, x_n, f)$ it can be seen that the pullback contains furthermore $|\beta|$ points, with the corresponding multiplicities given by the sequence β .

That is, we have for the pullback of the line L as divisor on C''

$$(f|_{C''})^*(L) = \sum_{i=1, \dots, m} \sum_{j=1, \dots, \alpha'_i} i \cdot x'_{i,j_i} + \sum_{i=1, \dots, m'} \sum_{j=1, \dots, \beta_i} i \cdot r'_{i,j_i} + \sum_{i=1, \dots, l} \sum_{j=1, \dots, \gamma_i} i \cdot t_{i,j_i},$$

where the set $\{t_{i,j_i}\} = C' \cap C''$ and the r'_{i,j_i} are general smooth points on C'' .

Let $\beta' := \beta + \gamma$, then $\beta' \geq \beta$. As the image $f|_{C''}(C'')$ is a curve of degree $d - 1$, we must have $I\alpha' + I\beta' = d - 1$.

C'' is a curve of genus $g' := g - |\gamma| + 1$.

The dimension of $V_n^{\Gamma, (\alpha, \beta)}(d, g) \cap \text{ev}_1^*(q)$ is equal to $2d + g + |\beta| - 1 + n - 2$ by theorem 3.43. With the same argument, the dimension of $V_{n'}^{\Gamma', (\alpha', \beta')}(d - 1, g')$ is

$$\begin{aligned} 2d - 2 + g' + |\beta'| - 1 + n' &= 2d - 2 + g - |\beta' - \beta| + 1 + |\beta'| - 1 + n' \\ &= 2d - 2 + g + |\beta| + n'. \end{aligned}$$

But the dimension must be equal, as we are only interested in components of $V_n^{\Gamma, (\alpha, \beta)}(d, g) \cap \text{ev}_1^*(q)$ of highest possible dimension — all other components are enumeratively irrelevant. Therefore we have $n' = n$, that is, the markings x_2, \dots, x_n must lie on C'' .

Let us determine one more condition the difference of the genera $g - g'$ needs to satisfy: the maximal number of points that can be in the intersection $C' \cap C''$ is $d - 1$ (because this is the degree of the image curve of C'' , and this image curve can not have more than $d - 1$ points in common with the line L .) Therefore $|\gamma| \leq d - 1$ and also $|\gamma| - 1 \leq d - 2$. But $|\gamma| - 1$ is equal to $g - g'$, so $g - g' \leq d - 2$.

With all these notations, the map

$$(C'', x'_{1,1}, \dots, x'_{1,\alpha'_1}, \dots, x'_{m,1}, \dots, x'_{m,\alpha'_m}, x_2, \dots, x_n, f|_{C''})$$

is an element of $V_{n-1}^{\Gamma', (\alpha', \beta')}(d - 1, g')$.

Therefore, there is a rational map from the component of K_2 containing this map to the stack $V_{n-1}^{\Gamma', (\alpha', \beta')}(d - 1, g')$ which forgets the component C' . This map is of degree $\binom{\beta'}{\beta}$, as there are $\gamma = \beta' - \beta$ possibilities how C' can be connected to C'' .

Summarizing, we have the following result proved in [32], Theorem 5.1:

3.46 Theorem

Let K be an irreducible component of $V_n^{\Gamma, (\alpha, \beta)}(d, g) \cap \text{ev}_1^*(q)$ (as defined in 3.42). Then

- $K \subset K_1$ is a component of

$$V_{n-1}^{\Gamma \cup \{q\}, (\alpha + e_k, \beta - e_k)}(d, g) \subset \overline{\mathcal{M}}'_{g, |\alpha| + n}(\mathbb{P}^2, d),$$

(where k denotes the contact order of $f(C)$ with the line L at the point q that is equal to one of the points r_{i,j_i}); or

- $K \subset K_2$ is a component that can be mapped with a degree $\binom{\beta'}{\beta}$ map to some

$$V_{n-1}^{\Gamma',(\alpha',\beta')}(d-1,g') \subset \overline{\mathcal{M}}'_{g',|\alpha'|+n-1}(\mathbb{P}^2,d-1),$$

where we have the following restrictions to the possible choices of α' , β' and g' :

$$\begin{aligned} \alpha' &\leq \alpha, \\ \beta' &\geq \beta, \\ I\alpha' + I\beta' &= d-1, \\ g - g' &= |\beta' - \beta| - 1 \text{ and} \\ d - 2 &\geq g - g'. \end{aligned}$$

All that remains to be done is to compute the multiplicities with which the different components occur in the intersection $V_n^{\Gamma,(\alpha,\beta)}(d,g) \cap \text{ev}_1^*(q)$. The results are stated in the following two propositions:

3.47 Proposition

In the intersection $V_n^{\Gamma,(\alpha,\beta)}(d,g) \cap \text{ev}_1^*(q)$, each component of K_1 (as in 3.46) belonging to $V_{n-1}^{\Gamma \cup \{q\},(\alpha+e_k,\beta-e_k)}(d,g)$ occurs with multiplicity k .

For a proof, see [32], Proposition 6.2.

3.48 Proposition

In the intersection $V_n^{\Gamma,(\alpha,\beta)}(d,g) \cap \text{ev}_1^*(q)$, each component of K_2 (as in 3.46) that can be mapped to $V_{n-1}^{\Gamma',(\alpha',\beta')}(d-1,g')$ occurs with multiplicity $I^{\beta'-\beta}$.

For a proof, see [32], Proposition 6.8.

3.49 Remark

The proofs of the propositions 3.47 and 3.48 are a very hard part of the proof of the Caporaso-Harris formula. One advantage of the tropical proof we will present in chapter 8 is that the analogous argument is a lot easier in the tropical setting.

Finally, let us sum up our results so far to derive the Caporaso-Harris formula. Consider once more the equality

$$N_{\text{cplx}}^{\alpha,\beta}(d,g) = (V_n^{\Gamma,(\alpha,\beta)}(d,g) \cap \text{ev}_1^*(q)) \cap (\text{ev}_2^*(p_2) \cap \dots \cap \text{ev}_n^*(p_n))$$

Here, we can now replace $(V_n^{\Gamma,(\alpha,\beta)}(d,g) \cap \text{ev}_1^*(q))$ by

$$\begin{aligned} &\sum_k k \cdot V_{n-1}^{\Gamma \cup \{q\},(\alpha+e_k,\beta-e_k)}(d,g) \\ &+ \sum \binom{\beta'}{\beta} \cdot \binom{\alpha}{\alpha'} \cdot I^{\beta'-\beta} \cdot V_{n-1}^{\Gamma',(\alpha',\beta')}(d-1,g'), \end{aligned}$$

where the second sum is taken over all β' , α' and g' satisfying the same requirements as needed for the Caporaso-Harris formula (see 3.37).

Intersecting this equation with $\text{ev}_2^*(p_2) \cap \dots \cap \text{ev}_n^*(p_n)$ gives the desired result, the formula of theorem 3.38.

3.50 Example

Let us revisit example 3.33. Now we want to apply the recursive relations given by the Caporaso-Harris formula to check again what we have done before, where we used only some pictures to convince ourselves. We want to determine $N_{\text{cplx}}^{(0),(3)}(3,0)$, that is, the number of rational cubics through 8 points. In the first step, we get

$$N_{\text{cplx}}^{(0),(3)}(3,0) = N_{\text{cplx}}^{(1),(2)}(3,0).$$

That is, we just get the same number if we count rational cubics that pass through one fixed point on the line L , and in addition through 7 points. The second step yields

$$N_{\text{cplx}}^{(1),(2)}(3,0) = N_{\text{cplx}}^{(2),(1)}(3,0) + N_{\text{cplx}}^{(0),(2)}(2,1),$$

where the second summand contributes 0, because there are no elliptic conics. That is, we can as well count the number of rational cubics that pass through 2 fixed points on L and in addition through 6 more points. This corresponds to our results from 3.33, where we noticed that the line L was just determined by the two points we specialized to lie on L .

The third step is the first one where we get some degenerations:

$$N_{\text{cplx}}^{(2),(1)}(3,0) = N_{\text{cplx}}^{(3),(0)}(3,0) + 2 \cdot N_{\text{cplx}}^{(0),(2)}(2,0) + 2 \cdot N_{\text{cplx}}^{(1),(1)}(2,1).$$

The last summand does not contribute, as before, because there are no elliptic conics. The factor 2 with which we count the second summand comes from the binomial coefficient $\binom{\beta'}{\beta}$, which counted the number of possibilities how the two components of the stable map can be attached to each other. This was precisely our argument before in 3.33. We can insert the recursively known result $N_{\text{cplx}}^{(0),(2)}(2,0) = 1$. In the last step finally, there is no contribution from the first type of component, K_1 :

$$N_{\text{cplx}}^{(3),(0)}(3,0) = 2 \cdot N_{\text{cplx}}^{(0),(0,1)}(2,0) + 3 \cdot N_{\text{cplx}}^{(1),(1)}(2,0) + N_{\text{cplx}}^{(0),(2)}(2,-1) + 3 \cdot N_{\text{cplx}}^{(2),(0)}(2,1).$$

The last summand does not contribute, as before. The first summand has to be counted with the factor 2. This time, the factor comes from the multiplicity $I^{\beta'-\beta}$, with which we have to count this component. The second summand comes with the factor 3, it is due to the possibilities $\binom{\alpha}{\alpha'}$ how the fixed points on L that the curve is required to meet can be distributed on the two components. We can insert now the recursively known terms and get $N_{\text{cplx}}^{(0),(0,1)}(2,0) = 2$, $N_{\text{cplx}}^{(1),(1)}(2,0) = 1$ and $N_{\text{cplx}}^{(0),(2)}(2,-1) = 3$, where the last equality is due to the possibilities how the markings can be arranged on the two lines.

Altogether, we have

$$\begin{aligned} N_{\text{cplx}}^{(0),(3)}(3,0) &= 2 \cdot N_{\text{cplx}}^{(0),(2)}(2,0) + 2 \cdot N_{\text{cplx}}^{(0),(0,1)}(2,0) + 3 \cdot N_{\text{cplx}}^{(1),(1)}(2,0) + N_{\text{cplx}}^{(0),(2)}(2,-1) \\ &= 2 + 4 + 3 + 3 = 12, \end{aligned}$$

so we get precisely what we had in 3.33 before.

3.51 Remark

In [4], section 1.4, Caporaso and Harris also showed an algorithm to count irreducible curves. The ideas to prove this formula are just the same, we only have to restrict to connected stable maps again, that is, we consider the substack corresponding to $V_n^{\Gamma,(\alpha,\beta)}(d,g)$ in $\overline{\mathcal{M}}_{g,|\alpha|+n}(\mathbb{P}^2,d)$, not in $\overline{\mathcal{M}}'_{g,|\alpha|+n}(\mathbb{P}^2,d)$.

As we will give a tropical proof for this formula in chapter 8, too, we shortly present this result here.

By $N_{\text{cplx}}^{\text{irr},(\alpha,\beta)}(d, g)$ we denote the numbers of irreducible curves satisfying the properties of definition 3.35 in addition.

They fulfill the following recursive relation:

$$\begin{aligned} N_{\text{cplx}}^{\text{irr},(\alpha,\beta)}(d, g) &= \sum_{k|\beta_k > 0} k \cdot N_{\text{cplx}}^{\text{irr},(\alpha+e_k, \beta-e_k)}(d, g) \\ &+ \sum \frac{1}{\sigma} \binom{2d+g+|\beta|-2}{2d_1+g_1+|\beta^1|-1, \dots, 2d_k+g_k+|\beta^k|-1} \\ &\cdot \binom{\alpha}{\alpha^1, \dots, \alpha^k} \\ &\cdot \prod_{j=1}^k \left(\binom{\beta^j}{\beta^j - \beta^{j'}} \cdot I^{\beta^{j'}} \cdot N_{\text{cplx}}^{\text{irr},(\alpha^j, \beta^j)}(d_j, g_j) \right) \end{aligned}$$

with the second sum taken over all collections of integers d_1, \dots, d_k and g_1, \dots, g_k and all collections of finite sequences $\alpha^1, \dots, \alpha^k$, β^1, \dots, β^k and $\beta^{1'}, \dots, \beta^{k'}$ satisfying

$$\begin{aligned} \alpha^1 + \dots + \alpha^k &\leq \alpha, \\ \beta^1 + \dots + \beta^k &= \beta + \beta^{1'} + \dots + \beta^{k'}, \\ |\beta^{j'}| &> 0, \\ d_1 + \dots + d_k &= d - 1 \text{ and} \\ g - (g_1 + \dots + g_k) &= |\beta^{1'} + \dots + \beta^{k'}| + k. \end{aligned}$$

Here as usual $\binom{n}{a_1, \dots, a_k}$ denotes the multinomial coefficient

$$\binom{n}{a_1, \dots, a_k} = \frac{n!}{a_1! \cdot \dots \cdot a_k! (n - a_1 - \dots - a_k)!}$$

and correspondingly, for sequences $\alpha, \alpha^1, \dots, \alpha^k$ the multinomial coefficient is

$$\binom{\alpha}{\alpha^1, \dots, \alpha^k} = \prod_i \binom{\alpha_i}{\alpha_i^1, \dots, \alpha_i^k}.$$

The number σ is defined as follows: Define an equivalence relation on the set $\{1, 2, \dots, k\}$ by $i \sim j$ if $d_i = d_j$, $g_i = g_j$, $\alpha^i = \alpha^j$, $\beta^i = \beta^j$ and $\beta^{i'} = \beta^{j'}$. Then σ is the product of the factorials of the cardinalities of the equivalence classes.

3.4. A SHORT OVERVIEW OF CURVES ON TORIC SURFACES

So far, we restricted our description of enumerative geometry to curves on \mathbb{P}^2 . The reason is that the main results which we cited here (and which we want to prove in the tropical context in the chapters 7 and 8) — Kontsevich's formula and the algorithm of Caporaso and Harris — are restricted to curves in \mathbb{P}^2 . However, there are also tropical curves which do not correspond to complex curves in \mathbb{P}^2 , but to complex curves on a different toric surface. Some results can be shown for these tropical curves, too. Therefore, we want to give a short overview based on [9] on toric surfaces and set up an enumerative problem

for curves on arbitrary toric surfaces, too. For more details about toric surfaces, we refer to [9].

3.52 Definition

The 2-dimensional *torus* is defined to be $T := (\mathbb{C}^*)^2$.

A *toric surface* is a normal variety X that contains T as a dense open subset, together with a *torus action* $T \times X \rightarrow X$ that extends the natural action of T on itself.

3.53 Example

\mathbb{P}^2 is a toric surface. Let \mathbb{P}^2 be given by the coordinates $\{x : y : z\}$. Then the torus T can be embedded as $(t_1, t_2) \mapsto (t_1 : t_2 : 1)$. The action

$$T \times \mathbb{P}^2 \rightarrow \mathbb{P}^2 : ((t_1, t_2), (x : y : z)) \mapsto (t_1 \cdot x : t_2 \cdot y : z)$$

extends the natural action of T .

The main idea about toric surfaces is that they can be described by a combinatorial structure — by a fan. This combinatorial structure helps to describe a lot of properties of the surface.


3.54 Definition

A *polyhedral cone* in \mathbb{R}^2 is a convex cone with apex at the origin, generated by finitely many vectors. A polyhedral cone is called *rational* if it is generated by lattice vectors, that is, by elements of \mathbb{Z}^2 . It is called *strongly convex* if it contains no line through the origin.

As all cones we are working with will have the described properties, we will by abuse of notation call a *strongly convex rational polyhedral cone* just a *cone*.

3.55 Example

Cones can be 2-dimensional, 1-dimensional, or 0-dimensional — the latter, if they just consist of the origin:



By a *face* of a cone we denote the cones that appear in the boundary — for a 2-dimensional cone, this can be the two half rays that limit σ or the origin itself; for a 1-dimensional cone, the origin is a face.

3.56 Definition

A *fan* \mathcal{F} is a set of cones such that


- each face of a cone $\sigma \in \mathcal{F}$ is contained in \mathcal{F} , too, and

- the intersection of two cones is a face of each.

A fan is called *complete* if it covers the whole of \mathbb{R}^2 .

3.57 Example

The following picture shows a complete fan:



It contains of 7 cones: 3 two dimensional cones, 3 one dimensional cones, and the origin.

3.58 Definition

By $M := \text{Hom}(\mathbb{Z}^2, \mathbb{Z})$ we denote the dual lattice of \mathbb{Z}^2 .

For a cone σ we define the *dual cone* σ^\vee to be the set of vectors in $M \otimes \mathbb{R}$ that are nonnegative on σ .

σ^\vee defines a semigroup $S_\sigma = \sigma^\vee \cap M$ which in turn defines a group algebra $\mathbb{C}[S_\sigma]$.

Such a group algebra defines an affine variety $U_\sigma := \text{Spec}(\mathbb{C}[S_\sigma])$, the *affine variety associated to σ* .

3.59 Example

Let σ be the origin. The dual cone σ^\vee is the whole plane. The intersection of σ^\vee with the dual lattice $M \cong \mathbb{Z}^2$ (hence, M itself) is as semigroup generated by the four vectors e_1 , e_2 , $-e_1$ and $-e_2$. We write the semigroup structure multiplicatively. The four generators of course fulfill two relations. We can therefore write the generators as x , x^{-1} , y and y^{-1} . Therefore, $\mathbb{C}[S_\sigma] = \mathbb{C}[x, x^{-1}, y, y^{-1}]$ and the affine variety associated to the origin, $U_\sigma := \text{Spec}(\mathbb{C}[S_\sigma])$, is just the torus.

3.60 Lemma

If τ is a face of a cone σ , then there is a map which embeds U_τ into U_σ as a principal open subset.

For a proof, see [9], section 1.3.

3.61 Definition

Let \mathcal{F} be a fan. Then we can define the *toric surface $X_{\mathcal{F}}$ associated to \mathcal{F}* in the following way: for all cones $\sigma \in \mathcal{F}$, we take the affine varieties U_σ as in 3.58 and glue them with the aid of the maps from lemma 3.60.

3.62 Remark

Note that $X_{\mathcal{F}}$ is indeed a toric surface: each fan contains the origin $(0,0)$ as a cone,


whose associated affine variety is T . The torus action is given on each U_σ by the inclusion $\mathbb{C}[S_\sigma] \rightarrow \mathbb{C}[S_\sigma] \otimes \mathbb{C}[S_{(0,0)}]$.

3.63 Remark

It is also possible to associate a fan to a given toric variety. In fact, it can be seen that the corresponding construction is an equivalence of categories.

3.64 Example

Let us compute the group algebras and varieties corresponding to the cones in the fan of example 3.57. The picture shows the three dual cones to the 2-dimensional cones.



These three dual cones are generated by $-e_1$ and $-e_2$, respectively e_1 and $(1, -1)$, respectively e_2 and $(-1, 1)$. In none of the cases fulfill the two generators any relations. We can therefore consider them as free generators of the semigroup S_σ . Hence each affine variety $U_\sigma = \text{Spec } \mathbb{C}[x, y]$ corresponding to the two dimensional cones is just a copy of \mathbb{C}^2 . By computing the gluing maps (which are given by the affine varieties associated to the 1-dimensional cones, the rays) we can see that $X_{\mathcal{F}} = \mathbb{P}^2$.

Here are some basic properties of toric surfaces that we can read off from the combinatorial description by the fan:

3.65 Lemma

The toric surface associated to a fan is compact if and only if the fan is complete.

3.66 Lemma

An affine toric variety U_σ is nonsingular if and only if there are generators of σ which can be completed to a lattice basis of \mathbb{Z}^2 .

For proofs, see [9].

We can deduce the (non-surprising) fact that \mathbb{P}^2 is a nonsingular compact surface.

Another important property of toric surfaces is that we can also interpret the divisor classes on a toric surface by means of the corresponding fan.

3.67 Lemma

Let \mathcal{F} be a complete fan with rays (that is, 1-dimensional cones) τ_1, \dots, τ_r . Then the divisor class group $\text{Div}(X_{\mathcal{F}})$ of $X_{\mathcal{F}}$ is generated by divisors D_{τ_i} defined by the rays τ_i , where D_{τ_i} is the closure of $\text{Spec } \mathbb{C}[\tau_i^\perp \cap M]$.

(τ_i^\perp denotes the set of all vectors which are orthogonal to τ_i .)

For a proof, see [9], section 3.3.

3.68 Remark

Let Δ be a lattice polytope in \mathbb{Z}^2 . We can draw a perpendicular ray to each side of Δ . The union of these rays defines a complete fan in \mathbb{R}^2 .

The integer length (see definition 2.11) of a side of Δ associates a number to the corresponding ray, and with this, to the corresponding divisor D_{τ_i} . Let s_1, \dots, s_r be the sides of the polygon, where each side s_i is dual to the ray τ_i . Let l_i denote the integer length of s_i . So the polygon defines the divisor class $l_1 \cdot D_{\tau_1} + \dots + l_r \cdot D_{\tau_r}$.

That is, the polytope defines not only a toric variety (via the fan) but also a divisor class in the toric surface.


Note that $\text{Div}(X_{\mathcal{F}})$ is not freely generated by the D_{τ_i} . The D_{τ_i} satisfy two relations which correspond to the condition that the polygon Δ closes up.

3.69 Definition

Let Δ_d be the triangle with vertices $(0, 0)$, $(d, 0)$ and $(0, d)$.

3.70 Example

The perpendicular lines of the triangle Δ_d from above obviously define the fan of \mathbb{P}^2 as in example 3.57. The integer lengths d of the sides define the divisor class of curves of degree d .



3.71 Definition

Due to remark 3.68 a polytope Δ defines a toric surface and a divisor class on it. We define $N_{\text{cplx}}(\Delta, g)$ (and $N_{\text{cplx}}^{\text{irr}}(\Delta, g)$) (analogously to definition 3.14 (and 3.18)) to be the number of (irreducible) curves of genus g and of the corresponding divisor class in the toric surface that pass through the appropriate number of points in general position.

Note that $N_{\text{cplx}}(\Delta_d, g) = N_{\text{cplx}}(d, g)$ with Δ_d as in 3.69 (and the same for the corresponding numbers of irreducible curves). As for the numbers $N_{\text{cplx}}(d, g)$, we can also interpret the more general numbers $N_{\text{cplx}}(\Delta, g)$ as intersection products on the moduli space of stable maps. The basic results from above hold in this case as well. Only, as mentioned before, both Kontsevich's formula and the algorithm of Caporaso and Harris are restricted to curves in \mathbb{P}^2 .

4. THE TROPICAL ENUMERATIVE PROBLEM IN THE PLANE

In this chapter, we would like to define tropical analogues $N_{\text{trop}}^{\text{irr}}(\Delta, g)$ of the numbers $N_{\text{cplx}}^{\text{irr}}(\Delta, g)$ (respectively, tropical analogues $N_{\text{trop}}(\Delta, g)$ of the numbers $N_{\text{cplx}}(\Delta, g)$) that were defined in 3.14 and 3.71. Also, we want to define tropical analogues of the moduli spaces of stable curves and stable maps. In chapter 3, we have seen that these moduli spaces were helpful to answer many classical enumerative problems. Of course, our aim is to make answers to enumerative problems easier by using tropical geometry — therefore, we would like to work with an easy, combinatorial description of tropical curves as mentioned in chapter 2. In the first two sections (4.1 and 4.2), we will introduce such a combinatorial description of a tropical curve, and define the analogues of the moduli spaces of stable curves, respectively stable maps. In the third section (4.3) we will compare the new definition of a tropical curve with the ones we made in chapter 2. We will study the moduli space of tropical curves in more detail in section 4.4.

An important result that we found in remark 3.16 for the numbers $N_{\text{cplx}}^{\text{irr}}(\Delta, g)$ (respectively, $N_{\text{cplx}}(\Delta, g)$ — see 3.18 and 3.71) was, that they do not depend on the position of the points through which we count the curves (as long as the points are sufficiently general). This statement was a consequence of the fact that the numbers $N_{\text{cplx}}^{\text{irr}}(\Delta, g)$ appear as intersection products on the moduli space of stable maps. In tropical geometry, we cannot work with intersection products, as intersection theory is not yet defined in its full extent in the tropical world. We therefore have to find another argument, why the numbers $N_{\text{trop}}^{\text{irr}}(\Delta, g)$ do not depend on the position of the points that the tropical curves are required to meet. We hope that our results will help to develop a theory of intersection products in the tropical world.

In the fifth section (4.5), we will make the enumerative problem precise and define the numbers $N_{\text{trop}}^{\text{irr}}(\Delta, g)$. The main theorem of this chapter — the theorem that the numbers $N_{\text{trop}}^{\text{irr}}(\Delta, g)$ do not depend on the position of the points — is stated in 4.53. For its proof, we need to equip the moduli space of tropical curves with more structure. In section 4.6 we will therefore introduce the structure of polyhedral complexes. In the last section (4.7) finally, we will prove theorem 4.53, using a local moduli space. The statement that also the numbers $N_{\text{trop}}(\Delta, g)$ do not depend on the position of the point will not be proved, the proof is analogous.

The main result of this chapter — theorem 4.53 — was achieved in joint work with Andreas Gathmann and published in [14] (though using different methods).

4.1. ABSTRACT TROPICAL CURVES

Following the ideas of chapter 2.3, we want to define parametrized tropical curves roughly as graphs in \mathbb{R}^2 that satisfy the balancing condition. But in this definition, we also want to take the ideas of chapter 3 into account. That is, we want to consider abstract tropical curves first (also abstract n -marked curves), and then maps from these abstract curves to \mathbb{R}^2 .

We will start with the definition of a graph.


4.1 Definition

Let I_1, \dots, I_k be closed (bounded or half bounded) real intervals. Choose some (not necessarily distinct) boundary points P_1, \dots, P_r and Q_1, \dots, Q_r of the intervals $I_1 \cup \dots \cup I_k$. The topological space Γ that is obtained by identifying P_i and Q_i for all $i = 1, \dots, r$ in $I_1 \cup \dots \cup I_k$ is called a *graph*. We fix the following notations:

- The boundary points of the intervals I_1, \dots, I_k are called the *flags* of Γ . The set of flags is denoted by Γ' .
- The images of the flags in Γ are called the *vertices* of Γ . If F is a flag its image vertex will be denoted by ∂F . The set of vertices is denoted by Γ^0 .
- For a vertex V we define the *valence* of V , $\text{val } V$, as the number of flags F with $\partial F = V$.
- The open intervals $I_1^\circ, \dots, I_k^\circ$ are open subsets of Γ , they are called the *edges* of Γ . The set of edges is denoted by Γ^1 . A flag F belongs to exactly one edge of Γ which will be denoted by $[F]$. We can therefore also think of a flag F as the edge $[F]$ together with a direction oriented away from the vertex ∂F .
- An edge is called *bounded* if its corresponding open interval is bounded, and *unbounded* otherwise. The set of bounded edges is denoted by Γ_0^1 , the set of unbounded edges by Γ_∞^1 . The unbounded edges will also be called *ends* of Γ .

4.2 Example

The following picture shows four unbounded and one bounded interval and a choice of boundary points p_i, q_i which are identified to get the graph Γ below.



We also need to define some properties of graphs:

4.3 Definition

A graph is called *connected* if Γ is connected as a topological space.

The *genus* of a graph is defined to be

$$g(\Gamma) := 1 - \#\Gamma^0 + \#\Gamma_0^1.$$

A graph is called *3-valent*, if for all vertices $V \in \Gamma^0$, we have $\text{val } V = 3$.

4.4 Remark

We can think of the graph Γ as a cell complex. Then we can see that

$$g(\Gamma) = 1 - \chi(\Gamma) = 1 - \dim H_0(\Gamma, \mathbb{Z}) + \dim H_1(\Gamma, \mathbb{Z}).$$

As $\dim H_0(\Gamma, \mathbb{Z})$ is equal to the number of connected components of Γ , a graph can only have negative genus if it is disconnected. So let us assume Γ has k connected components $\Gamma_1, \dots, \Gamma_k$. Then the genus $g(\Gamma_i)$ is nonnegative. As the sum of the number of vertices in each connected component is equal to $\#\Gamma^0$ and the same holds for $\#\Gamma_0^1$, we have

$$g(\Gamma_1) + \dots + g(\Gamma_k) = k - \#\Gamma^0 + \Gamma_0^1 = g(\Gamma) + k - 1.$$

If Γ is connected, the genus is equal to

$$g(\Gamma) = 1 - \dim H_0(\Gamma, \mathbb{Z}) + \dim H_1(\Gamma, \mathbb{Z}) = \dim H_1(\Gamma, \mathbb{Z}).$$

An element in $H_1(\Gamma, \mathbb{Z})$ is a collection of bounded edges that form a loop. The genus $g(\Gamma)$ is the number of “independent” loops of the graph Γ . (For further information on cell complexes and homology, see for example [31].)

We are now ready to define abstract marked tropical curves. Following an idea of Mikhalkin, markings should be thought of as unbounded edges (see [21]).

4.5 Definition


An *abstract tropical curve* is a graph Γ such that all vertices $V \in \Gamma^0$ have valence at least 3. An *n -marked abstract tropical curve* is a tuple $(\Gamma, x_1, \dots, x_n)$, where Γ is a tropical curve and $x_1, \dots, x_n \in \Gamma_\infty^1$ are distinct unbounded edges.

Two marked tropical curves $(\Gamma, x_1, \dots, x_n)$ and $(\Gamma', x'_1, \dots, x'_n)$ are isomorphic if there is a homeomorphism $\varphi : \Gamma \rightarrow \Gamma'$ mapping x_i to x'_i for all i , and such that every edge of Γ is mapped bijectively onto an edge of Γ' by an affine map of slope ± 1 , that is, by a map of the form $t \mapsto a \pm t$ for some $a \in \mathbb{R}$.

The set of all isomorphism classes of connected n -marked tropical curves with exactly n unbounded edges and of genus g is called $\mathcal{M}_{\text{trop}, g, n}$.

4.6 Example

There is exactly one tropical curve in $\mathcal{M}_{\text{trop}, 0, 3}$, since there is only one graph with 3 unbounded edges all of whose vertices have valence at least 3:



(we will always draw the marked unbounded edges as dotted lines). Hence $\mathcal{M}_{\text{trop}, 0, 3}$ is simply a point.

4.7 Remark

We can think of $\mathcal{M}_{\text{trop}, g, n}$ as a tropical analogue of the space of stable curves $\mathcal{M}_{g, n}$ as defined in 3.4. In fact, example 4.6 gives the expected result: both $\mathcal{M}_{\text{trop}, 0, 3}$ and $\mathcal{M}_{0, 3}$ are just points.


4.8 Remark

The isomorphism condition of definition 4.5 means that every edge of a marked tropical curve has a parametrization as an interval in \mathbb{R} that is unique up to translation and sign. In particular, every bounded edge e of a tropical curve has an intrinsic *length* that we will denote by $l(e) \in \mathbb{R}$. For an unbounded edge $e \in \Gamma_\infty^1$, we will say that $l(e) = \infty$.

For each flag, we can get rid of the ambiguity in translation and sign by choosing as a normal form the parametrization that maps the edge $[F]$ to the interval $[0, l([F])]$, where the flag F is mapped to 0. This parametrization is called the *canonical parametrization*.


4.9 Example

As a second example, let us consider $\mathcal{M}_{\text{trop}, 0,4}$. There are 4 different types of graphs with four unbounded ends all of whose vertices have valence at least 3:



The three types on the left differ from each other by the choice of two marked ends that come together at a 3-valent vertex. The fourth type has just one 4-valent vertex. For the first three types, the set of tropical curves of this type is parametrized by the length of the bounded edge that links the two 3-valent vertices. As the length is required to be bigger than or equal to 0, this set is parametrized by a half line. The fourth one does not have a bounded edge at all, therefore, it is parametrized by a point. In fact, it can be seen as the “limit” of all three other types. That is, the three half lines come together at a point which corresponds to a curve of type 4.

Hence, $\mathcal{M}_{\text{trop}, 0,4}$ looks like a tropical line:



In the picture, the type which is parametrized by each half line is drawn next to it.

Note that also here the theory coincides with the theory of stable maps, where $\overline{\mathcal{M}}_{0,4}$ is just a line (see 3.3).

4.2. PARAMETRIZED TROPICAL CURVES

We are now ready to define parametrized tropical curves:

4.10 Definition

A *parametrized tropical curve* is a pair (Γ, h) where Γ is an abstract tropical curve and $h : \Gamma \rightarrow \mathbb{R}^2$ is a continuous map such that:

- (1) On each edge e of Γ the map h is of the form

$$h(t) = a + t \cdot v$$

for some $a \in \mathbb{R}^2$ and $v \in \mathbb{Z}^2$ (i.e. “ h is affine linear with rational slope”). The integral vector v occurring in this equation if we pick for e the canonical parametrization with respect to a chosen flag F of e (see remark 4.8) will be denoted by $v(F)$ and called the *direction* of F .

- (2) At every vertex $V \in \Gamma^0$, the *balancing condition* is fulfilled:


$$\sum_{F \in \Gamma^1 | \partial F = V} v(F) = 0.$$

An *n -marked parametrized tropical curve* is a tuple $(\Gamma, h, x_1, \dots, x_n)$ where (Γ, h) is a parametrized tropical curve, and $x_1, \dots, x_n \in \Gamma_\infty^1$ are distinct unbounded edges of Γ that are mapped to a point in \mathbb{R}^2 by h (that is, $v(F) = 0$ for the corresponding flags).

Two n -marked tropical curves $(\Gamma, h, x_1, \dots, x_n)$ and $(\Gamma', h', x'_1, \dots, x'_n)$ are called *isomorphic* if there is an isomorphism $\varphi : (\Gamma, x_1, \dots, x_n) \rightarrow (\Gamma', x'_1, \dots, x'_n)$ of the underlying abstract curves as in definition 4.5 such that $h' \circ \varphi = h$.

4.11 Example

The following picture shows a parametrized tropical curve of genus 0:



The marked unbounded edges are, as always, drawn as dotted lines. Their images in \mathbb{R}^2 are just points. Note that the balancing condition is fulfilled at each vertex. Take for example the vertex to which x_1 is adjacent: x_1 is contracted to a point, that is $v(x_1) = 0$. The other two flags adjacent to that vertex have the directions $(-1, 0)$ and $(1, 0)$. Or take

the vertex on the right next to it: there are three flags that map to it, and their directions are $(-1, 0)$, $(0, -1)$ and $(1, 1)$.

4.12 Remark

The map h of a parametrized tropical curve $(\Gamma, h, x_1, \dots, x_n)$ does not need to be injective on the edges. It may happen that $v(F) = 0$ for a flag F , that is, the edge $[F]$ is contracted to a point in \mathbb{R}^2 . The remaining flags around the vertex ∂F then satisfy the balancing condition themselves. If ∂F is a 3-valent vertex, this means that the two other flags F_1 and F_2 around ∂F have to satisfy $v(F_1) = -v(F_2)$, that is, they point in opposite directions. Hence, the image $h(\Gamma)$ looks locally around $h(\partial F)$ like a straight line.

This holds in particular for the marked unbounded edges x_1, \dots, x_n , as they are required to be mapped to a point. Therefore, they can be seen as tropical analogues of the marked points of stable maps as in 3.8. By abuse of notation we will therefore often call the marked unbounded edges “*marked points*”.

Note that the contracted bounded edges also lead to “hidden moduli parameters”: if we vary the length of a contracted bounded edge, then we arrive at a family of different parametrized tropical curves whose images in \mathbb{R}^2 are all the same. This behavior is well-known for stable maps, too.

4.13 Remark

If the direction $v(F) \in \mathbb{Z}^2$ of a flag F of a plane tropical curve is not equal to zero then it can be written uniquely as a positive integer times a primitive integral vector, $v(F) = \omega(F) \cdot u(F)$. This positive integer $\omega(F)$ is what we called the *weight* of the corresponding edge in chapter 2.

4.14 Definition

The *degree* of an n -marked plane tropical curve is defined to be the unordered tuple

$$\Delta = (v(F); [F] \in \Gamma_\infty^1 \setminus \{x_1, \dots, x_n\})$$

of directions of its non-marked unbounded edges. If this degree consists of the vectors $(-1, 0)$, $(0, -1)$, $(1, 1)$ each d times then we simply say that the degree of the curve is d .

4.15 Remark

As in chapter 2.3, we can draw the dual to the image of a parametrized tropical curve. For each edge, we draw a line segment perpendicular to its image. We know that this dual is a subdivision of a polygon called Newton polygon.

In particular, we can draw the duals to the images of the unbounded edges $\Gamma_\infty^1 \setminus \{x_1, \dots, x_n\}$. This corresponds to the boundary of the Newton polygon. By abuse of notation, the Newton polygon is also denoted by Δ .

In chapter 6, we will see that parametrized tropical curves occur as limits of complex curves in toric surfaces. In 3.68 we have seen that a Newton polygon Δ defines a toric surface together with a divisor class. Also vice versa, a toric surface together with a complex curve determines a Newton polygon. We will see in chapter 6 that the parametrized tropical curves that occur as limits of a complex curve on a toric surface given by Δ are of degree

Δ . This justifies in particular why we call a tropical curve whose degree consists of the vectors $(-1, 0)$, $(0, -1)$, $(1, 1)$ each d times (or, in the dual language: whose degree is the triangle Δ_d with vertices $(0, 0)$, $(d, 0)$ and $(0, d)$, see definition 3.69) a tropical curve of degree d — they occur as limits of degree d complex curves in \mathbb{P}^2 .

Now we want to define a moduli space for parametrized tropical curves of a given degree.

4.16 Definition

For all $g, n \geq 0$ and Δ , let $\mathcal{M}_{\text{trop}, g, n}(\Delta)$ be the set of all isomorphism classes of connected parametrized tropical curves $(\Gamma, h, x_1, \dots, x_n)$ of degree Δ and genus $g' \leq g$.

4.17 Remark

As in the classical situation, we can also drop the requirement that Γ should be a connected graph. The space $\mathcal{M}'_{\text{trop}, g, n}(\Delta)$ of not necessarily connected parametrized tropical curves will help to determine the numbers $N_{\text{trop}}(\Delta, g)$ of not necessarily irreducible tropical curves. Here, we restrict to the connected case. Some arguments will be less complicated for connected tropical curves. However, the main ideas do not change when we allow disconnected curves, too.

We will study the space $\mathcal{M}_{\text{trop}, g, n}(\Delta)$ in more detail in chapter 4.4, but before we do so, let us compare the new definition of a tropical curve with the definitions from chapter 2.

4.3. THE COMPARISON OF PARAMETRIZED TROPICAL CURVES WITH THE DEFINITIONS OF TROPICAL CURVES FROM CHAPTER 2

From theorem 2.15 and remark 2.19 we know that each tropical curve associated to a complex curve over the completion of the field of Puiseux series (as defined in 2.4) (and with this, each tropical curve associated to a tropical polynomial as defined in 2.7) arises as the image of some parametrized tropical curve, because it looks like the image of a graph which fulfills the balancing condition. However, it is not self-evident that we can find a complex curve C for each parametrized tropical curve such that the image $h(\Gamma) \subset \mathbb{R}^2$ coincides with $\overline{\text{Val}(C)}$. (Respectively, it is not clear that we can find a tropical polynomial such that the tropical curve associated to it is equal to $h(\Gamma)$.)

In theorem 4.27, we will cite a result by Speyer that specifies which parametrized tropical curves actually come from a complex curve, that is, for which image $h(\Gamma) \subset \mathbb{R}^2$ of a parametrized tropical curve we have $h(\Gamma) = \overline{\text{Val} C}$ for a complex curve $C \subset (K^*)^2$. (Of course, there can be several parametrizations for a given image $h(\Gamma)$.)

But before we can do so, we have to define some more properties for parametrized tropical curves. It will not be possible to show in general that all parametrized tropical curves come from complex curves. In fact, it will only be possible for those parametrized tropical curves that satisfy a special property we are about to define.

4.18 Notation

Recall that we decided to work with the space $\mathcal{M}_{\text{trop}, g, n}(\Delta)$ (see 4.16 and 4.17). That is, we assume the underlying graphs Γ to be connected for the whole chapter, unless otherwise specified.

4.19 Definition


The *combinatorial type* of an abstract tropical curve is the homeomorphism class of Γ relative x_1, \dots, x_n , that is, the data of $(\Gamma, x_1, \dots, x_n)$ modulo homeomorphisms of Γ that map each marked point x_i to itself. (That is, the information about the length $l(e)$ of the edges of Γ is lost.)

The *combinatorial type* of a parametrized tropical curve $(\Gamma, h, x_1, \dots, x_n)$ is the data of the combinatorial type of the abstract tropical curve $(\Gamma, x_1, \dots, x_n)$ together with the direction $v(F)$ for each flags F of Γ' .

$\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ is defined to be the subset of $\mathcal{M}_{\text{trop}, g, n}(\Delta)$ of parametrized tropical curves of type α .

4.20 Example

The following picture shows a graph together with the information about the directions of all flags — hence a combinatorial type. (Note that for two flags F and F' which belong to the same edge $[F] = [F'] = e$, the direction $v(F)$ is equal to $-v(F')$. Therefore we only note one of these directions in our picture.) Below, the images of two tropical curves of this type are shown.



4.21 Lemma

For every combinatorial type α occurring in $\mathcal{M}_{\text{trop}, g, n}(\Delta)$ the space $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ is naturally an (unbounded) open convex polyhedron in a real vector space of dimension $2 + \#\Gamma_0^1$, that is a subset of a real vector space given by finitely many linear equations and finitely many linear strict inequalities.

Proof:

The combinatorial type fixes the graph Γ up to homeomorphism, and for each edge, the direction. It does not fix the length $l(e)$ of a bounded edge e of the graph. (Note that the

length of the image $h(e)$ is determined by the length $l(e)$ and the direction $v(F)$ of a flag F with $[F] = e$.) Nor does the combinatorial type fix the position of the image $h(\Gamma)$ in \mathbb{R}^2 .

Choose a “root vertex” $V \in \Gamma^0$. Two coordinates are given by the position of the image of that root vertex $h(V) \in \mathbb{R}^2$. The remaining coordinates are given by the lengths $l(e)$ of the bounded edges. That is, we can embed $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ into $\mathbb{R}^{2+\#\Gamma_0^1}$.

Of course, the lengths have to be nonnegative. Also, if the curve has higher genus, the lengths cannot be chosen independently. On the contrary, every loop leads to conditions on the lengths of the images of bounded edges in the loop. Therefore, the subset in $\mathbb{R}^{2+\#\Gamma_0^1}$ that corresponds to parametrized tropical curves is given by the linear inequalities that all lengths have to be nonnegative, and by the linear equations that the images of the loops have to close up in \mathbb{R}^2 . The polyhedron given by these conditions is unbounded, as for example the root vertex can be moved in the whole of \mathbb{R}^2 . \square

4.22 Remark

A different choice of the root vertex or of the order of the bounded edges in lemma 4.21 leads to a linear isomorphism on $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ of determinant ± 1 . This is obvious for the order of the bounded edges. If we choose another root vertex V' , the difference $h(V) - h(V')$ of the images of the two vertices is given by $\sum_F l([F]) \cdot v(F)$, where the sum is taken over a chain of flags leading from V to V' . This is obviously a linear combination of the lengths of the bounded edges, that is of the other coordinates of $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$. As these length coordinates themselves remain unchanged it is clear that the determinant of this change of coordinates is 1.

Note that not only $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ is embedded naturally into $\mathbb{R}^{2+\#\Gamma_0^1}$. If we fix a lattice in $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$, then this lattice is also uniquely embedded into $\mathbb{Z}^{2+\#\Gamma_0^1}$, as the determinant of a coordinate change is ± 1 . We will see later on that it seems natural to think of $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ as a vector space with a lattice (see example 4.57).

4.23 Remark

The dimension of the space $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ in $\mathbb{R}^{2+\#\Gamma_0^1}$ can be estimated. We know that $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ is a polyhedron given by some linear inequalities and some linear equalities. To estimate the dimension, we only have to check how many independent linear equalities there are. Let the genus of a tropical curve of type α be g' . Choose a set of g' generators of $H_1(\Gamma, \mathbb{Z})$. Such a generator is a loop in Γ , that is, it consists of some bounded edges that form a loop. The bounded edges are equipped with a direction by the fixed type α . The condition that the image of this loop closes up gives therefore 2 linear conditions on the lengths of the bounded edges in the loop. Altogether, we get $2g'$ linear conditions. As we chose generators of $H_1(\Gamma, \mathbb{Z})$, these conditions are in fact enough to guarantee that all loops close up. However, in general we do not know whether these conditions are independent (an example where this is not the case can be found in 4.34). For the dimension of $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$, we therefore get

$$\dim \mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta) \geq 2 + \#\Gamma_0^1 - 2g'.$$

4.24 Definition

The *expected dimension* of the space $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ (defined in 4.19), where α is a combinatorial type of a tropical curve of genus $g' \leq g$, is defined to be

$$\text{edim}(\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)) = 2 + \#\Gamma_0^1 - 2g'.$$

4.25 Remark

If Γ is a 3-valent graph of genus $g' \leq g$ with k unbounded edges, then it has $k - 3 + 3g'$ bounded edges. (If Γ has also higher valent vertices, then the number of bounded edges is smaller.) Therefore

$$\begin{aligned} \text{edim} \mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta) &\leq 2 + (n + \#\Delta) - 3 + 3g' - 2g' \\ &= n + \#\Delta - 1 + g' \\ &\leq n + \#\Delta + g - 1, \end{aligned}$$

with equality if and only if α is the combinatorial type of a 3-valent tropical curve of genus g . (The number of unbounded edges is prescribed with the degree Δ and the number of markings n .)

That is, we have to choose $n = \#\Delta + g - 1$ points in \mathbb{R}^2 in order to expect a finite number of tropical curves that meet them.

4.26 Definition

A parametrized tropical curve of combinatorial type α is called *regular*, if the dimension of the space $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ coincides with the expected dimension. Otherwise, it is called *superabundant*.

We are now ready to state the result of Speyer that tells us which parametrized tropical curves actually come from complex curves:

4.27 Theorem (Speyer's Theorem)

For every regular parametrized tropical curve (Γ, h) (without markings) there is a complex curve C defined over the completion K of the field of Puiseux series (as defined in section 2.1), such that $h(\Gamma)$ is equal to the tropical curve associated to C .

For a proof, see [27], Theorem 5.0.4.

4.28 Remark

Note that the image $h(\Gamma)$ of a superabundant parametrized tropical curve can allow a regular parametrization, too. This shows that the image $h(\Gamma)$ can come from a complex curve even though the chosen parametrization is superabundant. (For example, the images of the curves from example 4.34 with the double edges can also be parametrized by a graph where the double edges are replaced by one edge of higher weight. Such a parametrization is regular.) In fact, plane tropical curves can be parametrized regularly (see [23], corollary 2.24). Hence all plane tropical curves come from complex curves. As we work with parametrized tropical curves here however, we still make the difference between superabundant and regular parametrized curves.

4.29 Remark

Note the different notations in [27] and here. What we call “tropical curves associated to a curve C over the Puiseux series” is just called “tropical curve” by Speyer. What we call parametrized tropical curve is called “zero-tension curve” by Speyer. His notation puts more emphasis on what we are actually interested in. He wants to gain information about complex curves by using tropical curves, therefore, he only calls those objects a tropical curve that actually come from a complex curve. However, theorem 4.27 tells us that we consider at least in general cases the right objects — those who actually come from a complex curve — when considering parametrized tropical curves. We prefer to call the balanced graphs where we cannot be sure if they come indeed from a complex curve also (parametrized) tropical curves, because we need them as “boundary objects”. In order to avoid confusion, we always add “parametrized”, “associated to a complex curve” or “associated to a tropical polynomial” whenever it is necessary to specify what we are talking about.

4.4. THE MODULI SPACE OF PARAMETRIZED TROPICAL CURVES

Let us come back to the space $\mathcal{M}_{\text{trop}, g, n}(\Delta)$. In section 4.3, we defined for each combinatorial type α a subset of this space. We would now like to study how the space $\mathcal{M}_{\text{trop}, g, n}(\Delta)$ can be described using the subsets $\mathcal{M}_{\text{trop}, g, n}^{\alpha}(\Delta)$.

4.30 Lemma

For all g, n and Δ there are only finitely many combinatorial types occurring in the space $\mathcal{M}_{\text{trop}, g, n}(\Delta)$ (see definitions 4.16 and 4.19).

Proof:

For all $g' \leq g, n$ and Δ there are only finitely many homeomorphism classes of connected graphs Γ of genus g' and with $n + \#\Delta$ unbounded edges. Furthermore, by 2.15 the image $h(\Gamma)$ is dual to a subdivision of the polygon associated to Δ . In particular, this means that the absolute value of the entries of the vectors $v(F)$ is bounded in terms of the size of Δ . Therefore there are only finitely many choices for the direction vectors. \square

4.31 Proposition

Let α be a combinatorial type occurring in $\mathcal{M}_{\text{trop}, g, n}(\Delta)$. Then every point in $\overline{\mathcal{M}}_{\text{trop}, g, n}^{\alpha}(\Delta)$ (where the closure is taken in $\mathbb{R}^{2+\#\Gamma_0^1}$, see lemma 4.21) can naturally be thought of as an element in $\mathcal{M}_{\text{trop}, g, n}(\Delta)$. The corresponding map


$$i_{\alpha} : \overline{\mathcal{M}}_{\text{trop}, g, n}^{\alpha}(\Delta) \rightarrow \mathcal{M}_{\text{trop}, g, n}(\Delta)$$

maps the boundary $\partial\mathcal{M}_{\text{trop}, g, n}^{\alpha}(\Delta)$ to a union of strata $\mathcal{M}_{\text{trop}, g, n}^{\alpha'}(\Delta)$ such that α' is a combinatorial type with fewer internal edges than α . Moreover, the restriction of i_{α} to any inverse image of such a stratum $\mathcal{M}_{\text{trop}, g, n}^{\alpha'}(\Delta)$ is an affine map.

Proof:

Note that by the proof of 4.21 a point in the boundary of the open polyhedron $\mathcal{M}_{\text{trop}, g, n}^{\alpha}(\Delta) \subset \mathbb{R}^{2+\#\Gamma_0^1}$ corresponds to a tuple $(\Gamma, h, x_1, \dots, x_n)$ where some bounded edges e have length $l(e) = 0$. Such a curve is of a different combinatorial type then,

because the homeomorphism class of the graph has changed. For all edges e with length $l(e) = 0$ the two flags F and F' with $[F] = [F'] = e$ are identified. We can as well remove the edges of length 0 then. Note that the balancing condition will be fulfilled at the new vertices. Two examples what this can look like are shown in the following picture. The edges who tend to have length zero when we move towards the boundary of the open polyhedron $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ are drawn in bold.



Let Γ_1 be the graph which is obtained by removing the edges of length 0. Note that Γ_1 has fewer bounded edges than Γ . The tuple $(\Gamma_1, h|_{\Gamma_1}, x_1, \dots, x_n)$ is a parametrized tropical curve again, possibly of a smaller genus than $(\Gamma, h, x_1, \dots, x_n)$. This shows that the points in the boundary $\partial\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ can naturally be thought of as parametrized tropical curves in $\mathcal{M}_{\text{trop}, g, n}(\Delta)$ themselves. The combinatorial types α' that can occur in the boundary of $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$, that is, in the image $i_\alpha(\partial\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta))$, have by construction fewer bounded edges than α . Finally, it is clear that the restriction of i_α to the inverse image of any stratum $\mathcal{M}_{\text{trop}, g, n}^{\alpha'}(\Delta)$ is an affine map since the affine structure on any stratum is given by the position of the curve in the plane and the lengths of the bounded edges. \square

4.32 Definition

We will say that a type α' appears *in the boundary* of another type α , if there is a point in $\partial\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ that is identified with a curve of type α' (as in the proof of proposition 4.31).

4.33 Remark

Note that by proposition 4.31 it is possible that there is a type of genus $g' < g$ in the boundary of a type of genus g . This is the reason why we allowed tropical curves of a genus $g' < g$ in the space $\mathcal{M}_{\text{trop}, g, n}(\Delta)$: they occur as deformations of genus g tropical curves.


Let us consider the moduli space of tropical curves again, and recall that we want to use it to count the images $\overline{\text{Val}(C)}$ of curves C over the Puiseux series. Assume there is a curve C such that $\overline{\text{Val}(C)}$ is equal to the image $h(\Gamma)$ of a parametrized tropical curve. Then this property does not depend on the parametrization (Γ, h) of this image. Also, we have seen in 4.27 that we can only be sure that the image $h(\Gamma)$ comes from a complex curve if the parametrization (Γ, h) is not superabundant. For our purpose to count complex curves with the aid of tropical curves, the space $\mathcal{M}_{\text{trop}, g, n}(\Delta)$ therefore seems to be too big. We should exclude redundant parametrizations, that is, for example parametrizations with contracted bounded edges (that is, edges which are mapped to a point by h , see 4.12), and we should exclude superabundant curves. However, we cannot exclude all superabundant parametrized curves from the space $\mathcal{M}_{\text{trop}, g, n}(\Delta)$, as some can appear as deformations

of regular curves (as in the second case of example 4.34). That is, superabundant curves can lie in the boundary of a regular type. As we need a closed moduli space, we cannot exclude these superabundant curves.

4.34 Example

We are going to present two superabundant curves, one which is not in the boundary of a regular type, and one which is.


The following picture shows a 3-valent curve.



Vertices are drawn bold here to distinguish them between all the parallel edges. The number 2 occurring over two edges denotes their weights.

Such a “flat” loop as it is given by the double edge makes the tropical curve superabundant, as the two linear equations given by this loop are not independent. However this curve is not in the boundary of a regular curve.

As second example, assume now that a regular tropical curve $(\Gamma, h, x_1, \dots, x_n)$ contains the graph in the following picture on the left as a subgraph, and assume we move towards a point in the boundary of $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ where the edges drawn in bold tend to have length 0. This point corresponds to a tropical curve that has Γ_1 as subgraph, a graph with a double edge.



This curve has now again a flat loop, however, the loop does not end at edges which are also parallel to the loop.

Types with such a flat loop will play a role later one, therefore we make the following definition:

4.35 Definition

We call a type of a tropical curve of genus g where the graph contains Γ_1 as in example 4.34 as a subgraph and is 3-valent else an *exceptional type*.

Our aim is now to define a subset called relevant subset of $\mathcal{M}_{\text{trop}, g, n}(\Delta)$. It does not contain parametrizations with contracted bounded edges, and it contains only superabundant curves that appear as deformations of regular ones.

Note that all curves with a “flat” loop as in example 4.34 are superabundant. But only those can occur in the boundary of a regular curve, where the flat loop happens due to the

vanishing of some edges that were part of the loop. In particular, each vertex V adjacent to the flat loop must be at least 4-valent, and as such a vertex V comes from (at least) two 3-valent vertices, we will also have that the directions of the flags F with $\partial F = V$ span \mathbb{R}^2 . This motivates the following definition:

4.36 Definition

Let V be a vertex of a graph Γ such that there are two flags adjacent to V that point in the same direction — that is, for two flags F_1 and F_2 with $\partial F_1 = \partial F_2 = V$ we have $u(F_1) = u(F_2)$. ($u(F_i)$ denotes the primitive integer vector pointing in the direction of F_i as in remark 4.13, $u(F_i) \cdot \omega(F_i) = v(F_i)$). We call such a vertex a *vertex with a double edge*.

Let $C = (\Gamma, h, x_1, \dots, x_n)$ be a tropical curve. If C has no contracted bounded edges, and if for all vertices with double edges as above the directions of the flags adjacent to V span \mathbb{R}^2 , then C is called *relevant*. (In particular, every vertex with a double edge of C is at least 4-valent.)

(Note that the property of being relevant depends only on the combinatorial type of a tropical curve C , not on C itself. We will therefore also call the type of C relevant.)

We define $\widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta) \subset \mathcal{M}_{\text{trop}, g, n}(\Delta)$ to be the subset of relevant tropical curves which satisfy in addition the following property: if they are of genus $g' < g$, then they appear in the boundary of a type of genus g . $\widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta)$ is called the *relevant subset*.

4.37 Remark

Note that those 3-valent curves we exclude when passing to the relevant subset are not important for enumerative arguments: later on, we will see that we have to count them with multiplicity 0 (see definition 4.47).

Our next aim is to study the dimensions of the subsets $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ of $\widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta)$.

4.38 Definition

Let α be a combinatorial type of a relevant tropical curve C in $\widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta)$. We define its *codimension* to be

$$\text{codim } \alpha := \sum_{V \in \Gamma^0} (\text{val } V - 3) + g - g',$$

where $g' \leq g$ denotes the genus of C . Note that the codimension is always a nonnegative integer.


4.39 Remark

The idea of the concept of the codimension of a combinatorial type is that it should correspond to the codimension of the stratum $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ in the space $\widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta)$. Note that this is true for the “expected codimension”: the maximal expected dimension of a stratum is $2n$ as we have seen in remark 4.25. As a 3-valent graph has $n + \#\Delta - 3 + 3g'$ bounded edges, we can conclude that in general $\#\Gamma_0^1 = n + \#\Delta - 3 + 3g' - \sum(\text{val } V - 3)$. This leads to an expected dimension of $n + \#\Delta + g' - 1 - \sum(\text{val } V - 3) = 2n - \text{codim } \alpha$ for the type α . The following example shows that $\text{codim } \alpha$ is not always equal to the

codimension of the stratum $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ in $\widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta)$. (However, proposition 4.41 gives a lower bound for the codimension of $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ in $\widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta)$ at least.)

4.40 Example

Assume we consider curves of genus 3. Then the combinatorial type of the curve in the picture below has codimension 6 as there are two 4-valent and two 5-valent vertices. Therefore we would expect $\dim \mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta) = 2n - \text{codim } \alpha = 10 - 6 = 4$. We have $\dim \mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta) = 5$ however: there are two dimensions for translations of the curve in the plane, 2 for the bounded edges on the bottom which are not involved in any loop, and one more for rescaling the whole curve.



4.41 Proposition

Let $n = \#\Delta + g - 1$. For a relevant combinatorial type α occurring in $\widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta)$ (see 4.36) we have

$$\dim \mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta) \begin{cases} = 2n & \text{if } \text{codim } \alpha = 0; \\ = 2n - 1 & \text{if } \text{codim } \alpha = 1 \text{ or if } \alpha \text{ is exceptional}; \\ \leq 2n - 2 & \text{otherwise.} \end{cases}$$

Proof:

The proof of this proposition is based on the ideas of [23] proposition 2.23. Our result is similar to [26] lemma 2.2 but differs in that we consider parametrized tropical curves and not their images (so that we cannot apply Shustin's technique of Newton polygons).

By lemma 4.21 we know that $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ is the subset of $\mathbb{R}^{2+\#\Gamma_0^1}$ given by the conditions that all coordinates corresponding to lengths are positive, and that the loops close up in \mathbb{R}^2 . Let $g' \leq g$ denote the genus of a curve of type α .

By remark 4.23, we expect $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ to have dimension $2 + \#\Gamma_0^1 - 2g'$. If α is of codimension 0, then the expected dimension is $\text{edim } \mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta) = 2 + \#\Gamma_0^1 - 2g = 2 + (\#\Delta + n - 3 + 3g) - 2g = 2n$. If α is of codimension 1, the expected dimension is $2n - 1$. However, the $2g'$ equations given by the loops of Γ do not have to be independent. The expected dimension is only a lower bound. The aim is now to see that for relevant curves of codimension 0 or 1, the dimension is equal to the expected dimension, whereas for all other relevant curves, the codimension of $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ is at least 2 (with one exception: an exceptional type).

Pick a tropical curve $C = (\Gamma, h, x_1, \dots, x_n)$ in $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$. Choose a vertex V_1 of maximal valence. Let L be a line in \mathbb{R}^2 through $h(V_1)$. (For some cases later on, we choose a special line through $h(V_1)$, but for the moment, any line is good.)

Order the vertices of Γ starting with V_1 , such that the distance of their images in \mathbb{R}^2 from L is increasing and such that on every line parallel to L no vertex lies between two other with a lower number.


If there are more vertices besides V_1 on L , we first order the ones left of V_1 by increasing distance of V_1 , and then the ones right of V_1 . (For some cases later one, we will choose a different order for the points on the line L , but for the moment, this order is good.)

Orient the edges so that they point from the lower to the higher vertex. Unbounded edges are always oriented so that they point from their vertex to infinity.


Two properties of the chosen order and orientation are important for our reasoning:

- (1) at every vertex V , there is at least one edge that is oriented away from V and
- (2) if two edges whose images are parallel point to the same vertex, they point in the same direction.

The reason for the first property is the balancing condition.

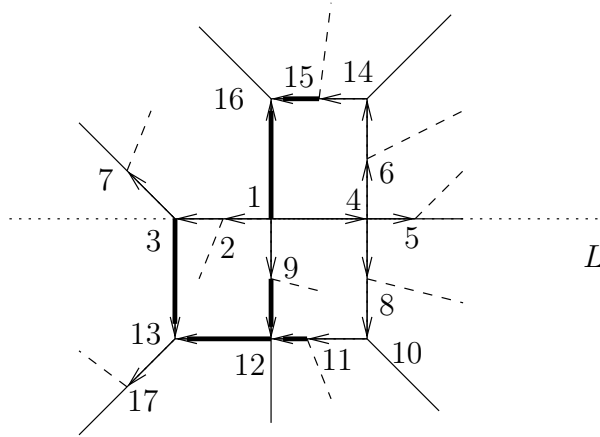


The second property means that two parallel edges that point to the same vertex cannot have opposite directions. Note first that two edges with opposite directions cannot point to the same vertex if their images are not parallel to L . So assume the two edges are mapped to a line parallel to L . Also, two edges that point to a vertex V have to be bounded. But this means that they end at another vertex which lies also on the line parallel to L . We have chosen the order of the vertices on a line parallel to L in such a way that at least one of the two other end vertices of the two bounded edges is of higher order than V . But this is a contradiction to the chosen orientation.




In the picture, the situation on the left is possible, the one on the right not.

We will now distinguish recursively $2g'$ edges $e_1, \dots, e_{g'}, e'_1, \dots, e'_{g'}$ as follows. For $i = 1, \dots, g'$ we let e_i be a (bounded) edge contained in a loop of $\Gamma \setminus \{e_1, \dots, e_{i-1}\}$ such that the vertex that this edge points to is maximal. Then $T := \Gamma \setminus \{e_1, \dots, e_{g'}\}$ is a maximal tree in Γ . In particular, for all $i = 1, \dots, g'$ the edge e_i closes a unique loop L_i in $T \cup \{e_i\} \subset \Gamma$. We let e'_i be the unique (bounded) edge of T that is contained in L_i and adjacent to the vertex that e_i points to.



The picture above shows an example. The edges e_i and e'_i are drawn in bold. Note that by construction the edges $e_1, \dots, e_{g'}$ are all distinct and different from the edges $e'_1, \dots, e'_{g'}$. It may happen however that not all edges $e'_1, \dots, e'_{g'}$ are distinct. An example is shown in the following picture:



In the picture, 4 is the highest number of a vertex contained in a loop. So we can for example choose the edge from 3 to 4 as e_1 . Then, if we remove this edge, 4 is still the highest number contained in a loop, and we can choose the edge from 2 to 4 as e_2 . Both edges e'_1 and e'_2 are equal to the edge from 1 to 4.

By construction, e_i and e'_i always point to the same vertex, namely to the highest vertex contained in the loop L_i .

We will now define a set of *conditional edges*. We start with the $2g'$ edges $e_1, \dots, e_{g'}, e'_1, \dots, e'_{g'}$ and remove some of these edges by applying the following rules at each vertex V :

- (i) if there is at least one edge e'_i pointing to V that is not parallel to its corresponding edge e_i then we keep the edge e'_i with this property such that i is maximal and remove all other edges $e'_1, \dots, e'_{g'}$ that point to V ;
- (ii) if there is no such edge then we remove all edges $e'_1, \dots, e'_{g'}$ that point to V .


The edges of $e_1, \dots, e_{g'}, e'_1, \dots, e'_{g'}$ which remain after this procedure are called *conditional edges*. Note that all edges $e_1, \dots, e_{g'}$ will end up to be conditional edges, and that all conditional edges will be distinct.

We claim that for any tropical curve in $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ the lengths of its conditional edges are determined uniquely in terms of the lengths of all other edges. To see this apply the following procedure recursively for $i = g', \dots, 1$: assume that we know already the


lengths of all edges $e_{i+1}, \dots, e_{g'}, e'_{i+1}, \dots, e'_{g'}$ as well as of all unconditional edges. Then by construction the only edges in the loop L_i whose lengths are not yet known can be e_i and e'_i (if V is the vertex that e_i and e'_i point to then all other edges in L_i must point to smaller vertices than V whereas all edges e_j and e'_j with $j < i$ point to vertices greater than or equal to V). If e_i and e'_i are not parallel then the condition that L_i closes up in \mathbb{R}^2 determines both their lengths uniquely. Otherwise e'_i is an unconditional edge by (ii), and e_i is again determined uniquely by the condition that L_i closes up.

It follows that the dimension of $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ is at most equal to $2 + \#\Gamma_0^1$ minus the number of conditional edges. So let us determine how many conditional edges there are. We claim that when passing from the $2g'$ edges $e_1, \dots, e_{g'}, e'_1, \dots, e'_{g'}$ to the conditional edges we removed at most $\text{val } V - 3$ edges at each vertex V .

To see this, let V be any vertex. In case (i) above there is at least one edge pointing away from V and one pair $\{e_i, e'_i\}$ that we do not remove.



In case (ii) all edges e'_i are parallel to e_i . By our chosen order, this is only possible if e'_i and e_i point in the same direction. But C is assumed to be a relevant curve, that is, the direction vectors of the flags adjacent to V have to span \mathbb{R}^2 . That is, we keep at least one e_i , at least one edge that is pointing away from V , and then there must be another edge adjacent to V , because otherwise \mathbb{R}^2 would not be spanned.



As all edges are pointing away from V_1 , we did not remove any edge at V_1 at all.

Therefore the number of conditional edges is at least

$$2g' - \sum_{V \neq V_1} (\text{val } V - 3).$$


Note that the number of bounded edges is equal to $\#\Gamma_0^1 = n + \#\Delta - 3 + 3g' - (\sum_V (\text{val } V - 3))$ by remark 4.25 and due to the fact that every vertex of higher valence can be separated by $(\text{val } V - 3)$ edges to get a 3-valent graph.

Therefore, an upper bound for the dimension of $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ is


$$\begin{aligned} & 2 + n + \#\Delta - 3 + 3g' - \left(\sum_V (\text{val } V - 3) \right) - 2g' + \sum_{V \neq V_1} (\text{val } V - 3) \\ &= n + \#\Delta + g' - 1 - (\text{val } V_1 - 3) \\ &= 2n - (g - g') - (\text{val } V_1 - 3). \end{aligned}$$

Now we consider several cases, stopping at the first one that applies to α :

- If $\text{codim } \alpha \leq 1$, then V_1 is the only vertex that can possibly have valence 4. So the number $2n - (g - g') - (\text{val } V_1 - 3) = 2n - \text{codim } \alpha$ and it follows that $\dim \mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta) = 2n - \text{codim } \alpha$ as we know already that this number is a lower bound by remark 4.23.
- If the number $2n - (g - g') - (\text{val } V_1 - 3)$ is at most $2n - 2$ then the statement follows immediately.
- If there are 2 vertices V_1 and V_2 of valence 4 such that there is no flat loop on the line connecting them between them as in example 4.34, we choose L to be the line through $h(V_1)$ and $h(V_2)$. As we already fixed the numbers 1 and 2 for these two vertices, we then cannot choose the order of the vertices on the line L as above. Instead, we choose the following order: We start with the vertices that lie between V_1 and V_2 , and order them correspondingly to their distance to V_1 . Next, we take the vertices on the opposite side of V_1 and order them, too, by their distance to V_1 . Then, we do the same thing with the vertices on the opposite side of V_2 :



The second property of our order (if two edges are parallel and point to a vertex, then they point in the same direction) is no longer fulfilled. In fact, there may be a vertex lying on the line L between V_1 and V_2 . Then if there are two parallel edges that point to V_r between V_1 and V_2 , they do not need to point in the same direction. In the picture above, we might for example have the following two edges oriented towards V_5 :



However, this property of the orientation was only important to see that we do not remove more than $\text{val } V - 3$ edges at every vertex V , when passing from the edges $e_1, \dots, e_{g'}$ to the conditional edges. Note that now we assumed that there is no flat loop between V_1 and V_2 . That is, there cannot be a loop such that the vertex V_r is the one with the highest number occurring in the loop. In particular, none of the edges parallel to L that point to V_r are contained in the set $e_1, \dots, e_{g'}$, and therefore, we do not remove any edge at V_r . With the same argument, we do not remove any edge at V_2 . Therefore, we can subtract $\text{val } V_2 - 3 \geq 1$ from the upper bound above and get $\dim \mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta) \leq 2n - 2$.

- The last case that remains to consider is the one of a type of genus g , where there are only 3 and 4-valent vertices, and all pairs of 4-valent vertices have a flat loop between them. In this case the dimension is obviously the same for the combinatorial type where each flat loop is replaced by one edge. This is then the type of a 3-valent curve of a lower genus. If there were more than one flat loop, it is a type of codimension at least 2 and we have already shown the statement for those. If there was only one flat loop, then the type was exceptional. The type where we replaced the loop by one edge is of codimension 1. Also for this type we have shown the statement already.

□

4.42 Corollary

Every 3-valent curve of genus g in $\widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta)$ is regular.

Proof:

By the above, we have seen that the dimension of $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ is equal to the expected dimension for a type with $\text{codim } \alpha = 0$. \square

Let us sum up the statements we made about the space $\widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta)$.

By lemma 4.30 and proposition 4.31 we can think of it as being obtained by starting with finitely many unbounded closed convex polyhedra $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ and then gluing them together by attaching each boundary $\partial\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ with affine maps to some polyhedra $\mathcal{M}_{\text{trop}, g, n}^{\alpha'}(\Delta)$. Proposition 4.41 gives us information about the dimensions of the strata. In particular, we know the highest occurring dimension, and we know the types α such that $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ is of codimension 1 in the relevant subset.

This structure (the “glued polyhedra”) will be called a *polyhedral complex* and we will study it in section 4.6. In fact, it would be nice if we could give the space $\widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta)$ the structure of an abstract tropical variety. However, the theory of abstract tropical varieties is very much in its beginnings, so we will not apply this language (see [22]). The structure of a polyhedral complex seems to be something that an abstract tropical variety should carry. Before we learn more about polyhedral complexes, we come back to tropical enumerative geometry.

4.5. THE TROPICAL ENUMERATIVE PROBLEM

Having defined the moduli space of relevant parametrized tropical curves, we want to come back to enumerative geometry.

4.43 Notation

The relevant subset $\widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta)$ (defined in 4.36) is a suitable moduli space for our following considerations. Therefore we will fix this space as moduli space of tropical curves for the whole chapter.

An important notion from chapter 3 are the evaluation maps. We will define their tropical analogues, too:

4.44 Definition

Let ev_i denote the i -th (tropical) evaluation map

$$\text{ev}_i : \widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta) \rightarrow \mathbb{R}^2 : (\Gamma, h, x_1, \dots, x_n) \mapsto h(x_i).$$

As x_i is contracted to a point in \mathbb{R}^2 by h , this is well-defined.

By ev we denote the product evaluation map

$$\text{ev} = \text{ev}_1 \times \dots \times \text{ev}_n : \widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta) \rightarrow \mathbb{R}^{2n} : (\Gamma, h, x_1, \dots, x_n) \mapsto (h(x_1), \dots, h(x_n)).$$

4.45 Lemma

On each subset $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$, the evaluation map ev_i from definition 4.44 is a linear map.

Proof:

The coordinates on $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ are by 4.21 given by a root vertex V and an order on the bounded edges. Of course, these coordinates do not need to be independent, but if they are not, they fulfill a linear condition themselves. As Γ is connected, we can reach x_i from the root vertex V by a chain of flags F , such that $[F]$ is a bounded edge. Then the position of $h(x_i)$ is given as a sum

$$h(V) + \sum_F v(F) \cdot l([F]),$$

where the summation goes over all flags F in the chain. Hence the position $h(x_i)$ is given by two linear expressions in the coordinates of $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$. \square

For a given set of points $\mathcal{P} = (p_1, \dots, p_n) \in \mathbb{R}^{2n}$, we would now like to count the tropical curves that pass through \mathcal{P} , that is, we would like to count inverse images of \mathcal{P} under ev . However, we have to count each inverse image with a certain multiplicity: As mentioned in the introduction, the main object of this thesis is to achieve enumerative results for complex curves with the help of tropical curves. More precisely, we want to count the limits of amoebas $\text{Log}(C)$ of complex curves C (respectively, images $\overline{\text{Val}(C)}$) instead of the complex curves themselves (see chapter 2). Each tropical curve comes therefore with a natural *multiplicity*: there might be several complex curves C_1, \dots, C_k such that the limits of their amoebas are all equal to the same tropical curve. Then we have to count this tropical curve k times.

This is only the idea why tropical curves have to be counted with a multiplicity at all. We will not define the multiplicity in that way. Instead, we will define it in some combinatorial way, and give an idea in chapter 6 why this definition coincides with the number of complex curves that map to the tropical curve (under Log and the limiting process which we will specify in chapter 6).

We need another definition before we can define the multiplicity:

4.46 Definition

Let $C = (\Gamma, h, x_1, \dots, x_n)$ be a 3-valent parametrized tropical curve.

- (1) A *string* in C is a subgraph of Γ homeomorphic either to \mathbb{R} or to S^1 (that is, a “path” starting and ending with an unbounded edge, or a path around a loop) that does not intersect the closures $\overline{x_i}$ of the marked points.
- (2) A tropical curve is called *rigid* if it contains no strings.

4.47 Definition

Let $C = (\Gamma, h, x_1, \dots, x_n)$ be a 3-valent parametrized tropical curve of degree Δ and genus g' , and let $n = \#\Delta + g' - 1$.

- (1) The *multiplicity of a vertex* $V \in \Gamma^0$ is defined to be the absolute value of the determinant $\det(v_1, v_2)$, where v_1 and v_2 are two directions of flags adjacent to V . The balancing condition tells us that it makes no difference which two of the three flags adjacent to V we choose.

- (2) If C is rigid and regular, the *multiplicity* of C is defined to be the product of the multiplicities of all vertices that are not adjacent to a marked point. Otherwise, the multiplicity of C is defined to be 0.

Note that the multiplicity of a tropical curve $C = (\Gamma, h, x_1, \dots, x_n)$ depends only on the type α , not on the lengths of the bounded edges.

As already mentioned, we will see that for a rigid and regular tropical curve, the multiplicity is equal to the number of complex curves of the corresponding degree and genus that map to it (under Log and the limiting process). A reason why we define the multiplicity of a nonrigid or superabundant tropical curve to be 0 is given in proposition 4.49.

Note that this definition of multiplicity differs from the definition given in [23], definition 4.15. There the multiplicity of a nonrigid or superabundant curve is not set to be 0.

4.48 Remark

Note that the multiplicity of a 3-valent nonrelevant curve is 0: either it contains a contracted bounded edge, or it contains a vertex where the edges do not span \mathbb{R}^2 , both leading to a vertex of multiplicity 0.

4.49 Proposition

Let C be a 3-valent curve of genus g and type α in $\widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta)$, and let $n = \#\Delta + g - 1$. Then $\text{mult } C = 0$ if and only if the evaluation map restricted to $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ is not injective.

Proof:


Let $\text{mult } C = 0$. We will study different cases:

- C is not rigid,
- C is not regular,
- there is a vertex of multiplicity 0.

First we note that the second case cannot occur, as every relevant 3-valent curve of genus g in $\widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta)$ is regular by 4.42.

Next, we deal with the third case. Assume V is a vertex of multiplicity 0. Then the flags adjacent to V do not span \mathbb{R}^2 . As C is relevant, it does not have a contracted bounded edge, hence no edge of direction 0. As C is 3-valent, we have 3 flags adjacent to V , and as none is of direction 0, two have to point in the same direction. But this is a contradiction to the relevance of C again, as the edges around V need to span \mathbb{R}^2 in this case.

It remains to deal with the first case — C is not rigid. Then $C = (\Gamma, h, x_1, \dots, x_n)$ contains a string $\Gamma_1 \subset \Gamma$. There is a deformation of C that changes the position of the string, but neither the images of the marked points nor the lines, on which the images of the edges in $\Gamma \setminus \Gamma_1$ lie. In particular, we can see that if we fix the position of the images of the marked points, then there is a whole family of curves of the same combinatorial type such that the marked points are mapped to the fixed images. Therefore $\text{ev} |_{\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)}$ is not injective. The following picture shows two examples of (images of) non rigid tropical curves, and their deformations:



For the other direction, assume ev is not injective on $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$. The set $(\text{ev} |_{\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)})^{-1}(\mathcal{P})$ is the set of all curves of type α that pass through \mathcal{P} . Let C be such a curve. Then $C + \ker(\text{ev} |_{\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)}) = (\text{ev} |_{\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)})^{-1}(\mathcal{P})$. Let $v \neq 0$ be a vector in $\ker(\text{ev} |_{\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)})$. Then v cannot only have nonzero entries at the coordinates corresponding to the root vertex, as a translation of C , $C + \lambda \cdot v$ will not pass through \mathcal{P} . Therefore v has at least one nonzero entry at a coordinate corresponding to the length of a bounded edge. This bounded edge is not mapped to a point, as C is relevant. But then also the image of $C + \lambda \cdot v$ in \mathbb{R}^2 contains a longer (respectively shorter) bounded edge, and is still passing through \mathcal{P} .

For an edge e (respectively, for a vertex V), we will say that *the position of e* (respectively V) is *fixed* in the family $C + \lambda v$, if the line (point) to which e (respectively V) is mapped is the same for all curves in the family. Otherwise, we say that an edge (vertex) *changes position*. Note that a vertex adjacent to a marked point cannot change position, as the marked points are mapped to the same image points \mathcal{P} in the family.

Remember also that at each vertex V of C the edges e_1 , e_2 and e_3 adjacent to V span \mathbb{R}^2 . (We assumed that no edge is mapped to a point, that is, if e_1 , e_2 and e_3 do not span \mathbb{R}^2 , two edges point in the same direction. But this cannot happen due to the relevance.) Therefore the positions of two of the edges e_i determine the position of V . If at least two edges do not change position, then also V does not change position. Hence at each vertex V that changes position there are at least two edges adjacent to V which change position.

We know that in the family $C + \lambda v$ there is an edge whose image grows longer (respectively shorter). Therefore there must also be a vertex V which changes position.

Starting at V , we can follow one, say e_1 , of the (at least) two edges adjacent to V that change position to the next vertex V_2 (so V_2 is the second vertex adjacent to e_1). V_2 must change position, too. Therefore, there are again at least two edges adjacent to V_2 which change position, e_1 , and another, e_2 . So we can follow e_2 to the next vertex and so on, until we either reach our vertex V again, or until we reach an unbounded end, which is not mapped to a point. In the second case, we can follow the second edge adjacent to V which changes position, until we either reach V again or another unbounded end. In any case, we finally end up with either a loop or a chain connecting two unbounded edges that change position. We cannot have met a marked point in this procedure, as the marked points do not change position. Therefore this loop respectively this chain connecting two unbounded edges is a string. Hence C is not rigid. \square

4.50 Lemma

Let $C = (\Gamma, h, x_1, \dots, x_n)$ be a 3-valent curve of genus g .

Assume $n = \#\Delta + g - 1$. Then C is rigid if and only if every connected component of $\Gamma \setminus \bigcup_i \overline{x_i}$ has genus 0 and exactly one unbounded edge.

Assume $n < \#\Delta + g - 1$. Then C is not rigid.

Proof:

Let $n = \#\Delta + g - 1$. Let us first assume that every connected component of $\Gamma \setminus \bigcup_i \overline{x_i}$ has genus 0 and exactly one unbounded edge. Then C contains obviously no string and is therefore rigid. Now assume C is rigid. Then every connected component of $\Gamma \setminus \bigcup_i \overline{x_i}$ needs to be rational and cannot have more than one unbounded edge. It remains to see that there cannot be a component with no unbounded edge. To see this, remove the closures of the marked points $\overline{x_1}, \dots, \overline{x_n}$ from Γ one after the other. As C is 3-valent, each removal can either separate one more component, or break a loop. As we end up with rational components, g of the removals must have broken a loop. That is, we have $1 + n - g = \#\Delta$ connected components. Therefore, each component has precisely one unbounded edge.

Now let $n < \#\Delta + g - 1$. Consider again $\Gamma \setminus \bigcup_i \overline{x_i}$. If one of the connected components is not rational, then C contains a string. So let us assume all connected components are rational. With the same argument as above, we can then see that there are $\#\Delta - 1$ connected components. Therefore, there must be at least one which contains more than one unbounded edge. Hence C contains a string. \square

We would like to count tropical curves whose marked points are mapped to a certain given set of points $\mathcal{P} = (p_1, \dots, p_n) \in \mathbb{R}^{2n}$. That is, we would like to count the tropical curves in $\text{ev}^{-1}(\mathcal{P})$. We already said that we have to count the tropical curves with their multiplicity. However, the multiplicity is only defined for 3-valent curves. (Recall that 3-valent relevant curves are regular by 4.42, therefore they are images of a complex curve by theorem 4.27.) Therefore we have to make a restriction on the configurations of points — they have to be in (tropical) general position.

4.51 Definition

A set of points $\mathcal{P} = (p_1, \dots, p_n) \in \mathbb{R}^{2n}$ is called to be in *(tropical) general position*, if the types of all tropical curves $C \in \text{ev}^{-1}(\mathcal{P})$ are of codimension 0 as defined in 4.38, that is, C is 3-valent and of genus g .

Note that the subset $U \subset \mathbb{R}^{2n}$ of points in general position is the complement of a union of polyhedra of dimension less than $2n$. This is true because due to proposition 4.41 the strata of types which are not of codimension 0 are of dimension less than $2n$, and they are mapped linearly to \mathbb{R}^{2n} by ev . It is in fact a consequence of our main theorem, 4.53, that U is not only open and of top dimension, but even dense in \mathbb{R}^{2n} .

We are now ready to formulate the *tropical enumerative problem*:

4.52 Definition

For all g, n and Δ , and for a set $\mathcal{P} = (p_1, \dots, p_n) \in \mathbb{R}^{2n}$ of points in general position as

defined in 4.51 (where $n = \#\Delta + g - 1$), we define *the number of tropical curves through \mathcal{P} counted with multiplicity* as

$$N_{\text{trop}}^{\text{irr}}(\Delta, g, \mathcal{P}) = \sum_{C \in \text{ev}^{-1}(\mathcal{P})} \text{mult}(C),$$

where ev is defined in 4.44, and mult in 4.47.

Note that the sum on the right is indeed finite: by lemma 4.30, there are only finitely many types α occurring in the space $\widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta)$. If there is a type α such that $\mathcal{M}_{\text{trop}, g, n}^{\alpha}(\Delta)$ contains infinitely many preimages $C \in \text{ev}^{-1}(\mathcal{P})$, then the multiplicity of these curves is 0 due to proposition 4.49.

The aim is to show that the number $N_{\text{trop}}^{\text{irr}}(\Delta, g, \mathcal{P})$ does not depend on \mathcal{P} . (Therefore, we will also denote it by $N_{\text{trop}}^{\text{irr}}(\Delta, g)$ in the following.) More precisely, we want to prove the following theorem:

4.53 Theorem

Let $U \subset \mathbb{R}^{2n}$ denote the subset of points in general position. Then the map

$$U \rightarrow \mathbb{N} : \mathcal{P} \mapsto N_{\text{trop}}^{\text{irr}}(\Delta, g, \mathcal{P})$$

is constant (where $N_{\text{trop}}^{\text{irr}}(\Delta, g, \mathcal{P})$ is defined in 4.52).

We could prove theorem 4.53 with the aid of Mikhalkin's Correspondence Theorem (see theorem 6.1, respectively [23], theorem 1). Roughly, it states that the number of tropical curves through a set of points $\mathcal{P} = (p_1, \dots, p_n) \in \mathbb{R}^{2n}$ counted with multiplicity coincides with the number of complex curves in the toric surface defined by Δ through a set of points (q_1, \dots, q_n) with $p_i = \log |q_i|$. As the latter number $N_{\text{cplx}}^{\text{irr}}(\Delta, g)$ does not depend on the position of the points (see 3.16 and 3.71), the result follows.

However, in this chapter we want to give a proof of theorem 4.53 within tropical geometry which imitates the methods of stable maps.

Before we can start with the proof, we need to equip the moduli space of relevant tropical curves with more structure.

4.6. THE MODULI SPACE AND THE STRUCTURE OF POLYHEDRAL COMPLEXES

4.54 Definition

Let X_1, \dots, X_N be (possibly unbounded) open convex polyhedra in real vector spaces. A *polyhedral complex* with cells X_1, \dots, X_N is a topological space X together with continuous inclusion maps $i_k : \overline{X_k} \rightarrow X$ such that X is the disjoint union of the sets $i_k(X_k)$ and the "coordinate changing maps" $i_k^{-1} \circ i_l$ are linear (where defined) for all $k \neq l$. We will usually drop the inclusion maps i_k in the notation and say that the cells X_k are contained in X .

The *dimension* $\dim X$ of a polyhedral complex X is the maximum of the dimensions of its cells. We say that X is of *pure dimension* $\dim X$ if every cell is contained in the closure of a cell of dimension $\dim X$. A point of X is said to be *in general position* if it is contained in a cell of dimension $\dim X$.

4.55 Definition

A *morphism between two polyhedral complexes* X and Y is a continuous map $f : X \rightarrow Y$ such that for each cell $X_i \subset X$ the image $f(X_i)$ is contained in only one cell of Y , and $f|_{X_i}$ is a linear map (of polyhedra).

Let X and Y be two polyhedral complexes of the *same* pure dimension, and let $f : X \rightarrow Y$ be a morphism. Let $P \in X$ be in general position, and such that also $f(P) \in Y$ is in general position. Then locally around P the map f is a linear map between vector spaces of the same dimension. We define the *multiplicity of f at P* $\text{mult}_f(P)$ to be the absolute value of the determinant of this linear map. Note that the multiplicity depends only on the cell of X in which P lies. We will therefore also call it the multiplicity of f in this cell.

Now let $Q \in Y$ be a point in general position. It is defined to be *in f -general position*, if all points $P \in f^{-1}(Q)$ are in general position in X . Note that the set of points in f -general position in Y is the complement of a subset of Y of dimension at most $\dim Y - 1$, therefore it is a dense open subset.

For a point $Q \in Y$ in f -general position we define the *degree of f at Q* to be

$$\deg_f(Q) = \sum_{P \in f^{-1}(Q)} \text{mult}_f(P).$$

This sum is indeed finite: first of all there are only finitely many cells in X . Moreover, in each cell (of maximal dimension) of X where f is not injective (that is, where there might be infinitely many inverse image points of Q) the determinant of f is zero and with it also the multiplicity for all points in this cell.

Note that the definition of the multiplicity $\text{mult}_f(P)$ in general depends on the coordinates we choose for the cells. However, we will use this definition only for a morphism for which the determinant does not depend on the chosen coordinates, if they are chosen in a natural way.

4.56 Lemma

The space $\widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta)$ (defined in 4.36) is a polyhedral complex of pure dimension $2n$ and $\text{ev} : \widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta) \rightarrow \mathbb{R}^{2n}$ a morphism of polyhedral complexes of the same dimension.

Proof:

The cells are obviously the sets $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ corresponding to relevant types. By 4.21 they are open convex polyhedra. By proposition 4.31, their boundary is also contained in $\widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta)$, and the coordinate changing maps are linear. By proposition 4.41, the highest dimension of such a cell is $2n$. Furthermore, by definition of the relevant subset each type of a lower genus is contained in the boundary of a type of genus g . Each higher-valent vertex can be resolved to 3-valent vertices. Therefore each type is contained in the boundary of a type of codimension 0. The evaluation map maps every cell of $\widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta)$ obviously to the only cell of the polyhedral complex (of pure dimension $2n$) \mathbb{R}^{2n} . Furthermore, ev restricted to a cell is a linear map due to 4.45. Therefore ev is a morphism between polyhedral complexes of the same dimension. \square

In fact, it would be desirable to see that $\text{mult}_{\text{ev}}(C)$ is equal to the multiplicity $\text{mult } C$ of a tropical curve C . Then a different formulation of theorem 4.53 would be to show that the map


$$\mathcal{P} \mapsto \text{deg}_{\text{ev}}(\mathcal{P})$$

is constant on U . However, the computation of $\text{mult}_{\text{ev}}(C)$ requires the knowledge of a basis of the space $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$. That is, we would need to find a set of vectors that span the vector space given by the conditions of the loops. It is of course possible to compute such a set of vectors for special given conditions, but not in general.

An example shows furthermore that in general the ev-multiplicity $\text{mult}_{\text{ev}}(C)$ will not be equal to the multiplicity of C :

4.57 Example

The picture shows a curve of genus 1. Before we can calculate $\text{mult}_{\text{ev}}(C)$ we have to find a basis of the parametrizing space of curves of this type. The root vertex is chosen to be V as in the picture, and we have six coordinates for the lengths of the bounded edges. The directions are indicated in the picture.



The two conditions of the loop are

$$1 \cdot l_2 - 1 \cdot l_4 = 0 \quad \text{and} \quad 1 \cdot l_2 + 1 \cdot l_4 - 2 \cdot l_5 - 2 \cdot l_6 = 0.$$

That is, a basis that generates the subspace in \mathbb{R}^8 defined by the two equations is for example given by the following 6 vectors:

$$e_1, e_2, e_3, e_5, (0, 0, 0, 1, 0, 1, 1, 0) \text{ and } (0, 0, 0, 0, 0, 0, 1, -1),$$

where the coordinates are chosen in the following order: first the two coordinates given by the root vertex, then the coordinates of the lengths as indicated in the picture. Also, we have to choose a chain of flags from V to each of the marked points x_i , in order to describe the position of $h(x_i)$ depending on the lengths of the edges in the chain. Let us

choose the following:

$$\begin{aligned} h(x_1) &= h(V) + l_1 \cdot \begin{pmatrix} -3 \\ -1 \end{pmatrix}, \\ h(x_2) &= h(V) + l_2 \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix} + l_3 \cdot \begin{pmatrix} 0 \\ 2 \end{pmatrix}, \text{ and} \\ h(x_3) &= h(V) - l_6 \cdot \begin{pmatrix} -2 \\ 0 \end{pmatrix}. \end{aligned}$$

If we evaluate this map at our chosen basis vectors, we get the following matrix as representation of ev :

$$\begin{pmatrix} 1 & 0 & -3 & 0 & 0 & 0 \\ 0 & 1 & -1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 2 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 2 \\ 0 & 1 & 0 & 0 & 0 & 0 \end{pmatrix}$$

The determinant of this matrix can be computed to be 4, but the multiplicity of the curve is 8. The reason for this is that we forget too many things that are important for the structure of a tropical variety. Without having defined what a tropical variety should be, we can say that it should also be equipped with a lattice (see remark 4.22) — in the case of a tropical curve the integer lattice decides about the weight that is associated to an edge of a certain direction. The space $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ in the above example is given as the intersection of two subspaces defined by the two equations that the loop closes. So we also have to take into account how the two lattices of these two subspaces intersect. Also, the coefficients of the equations of the loop can have a greatest common divisor which is bigger than 1. Then the subspace defined by the equations is the same if we divide out this factor. But the multiplicity of the curve is different if this factor is missing. In the above example, the index of the two lattices of the two subspaces each given by one equation can be computed to be 2 — the missing factor we needed to get the multiplicity of the curve. In fact, it can be seen for a set of examples (of genus 1) that this is the correct solution: the product of the determinant of ev , the greatest common divisors of the coefficients in the equations of the loop and the index of the two lattices of the two subspaces each given by one equation is equal to the multiplicity of a curve. However, this product is difficult to compute in general, as we need to compute lattice bases for subspaces given by arbitrary equations. For more information on lattices, we refer to [20].

As $\text{mult}_{\text{ev}}(C)$ is not equal to $\text{mult } C$ in general, we intend to define another morphism, which imitates the evaluation map in some sense, but has the advantage that its multiplicity at C coincides with $\text{mult } C$. For this, we will in fact need a different space, which is bigger than our moduli space of tropical curves. However, this definition will only work locally. Before we start with the definition of this local bigger moduli space, we collect the other ingredients which will be needed for the proof of theorem 4.53.

4.7. THE PROOF OF THEOREM 4.53

Theorem 4.53 states that the map $\mathcal{P} \mapsto N_{\text{trop}}^{\text{irr}}(\Delta, g, \mathcal{P})$ is constant. We want to prove this locally.

4.58 Lemma

The map $\mathcal{P} \mapsto N_{\text{trop}}^{\text{irr}}(\Delta, g, \mathcal{P})$ is locally constant at a point configuration \mathcal{P}' in the subset of \mathbb{R}^{2n} of points in general position.

Proof:

If \mathcal{P}' is a configuration of points in general position then at any curve that counts for $N_{\text{trop}}^{\text{irr}}(\Delta, g, \mathcal{P}')$ with a non-zero multiplicity the map ev is a local isomorphism: let $C \in \text{ev}^{-1}(\mathcal{P}')$ be a curve of type α with a nonzero multiplicity. By proposition 4.49, we know that $\text{ev}|_{\mathcal{M}_{\text{trop}, g, n}^{\alpha}(\Delta)}$ is injective. Because \mathcal{P}' is in general position, we know that $\text{codim } \alpha = 0$ and therefore the dimension of $\mathcal{M}_{\text{trop}, g, n}^{\alpha}(\Delta)$ is $2n$ by proposition 4.41. Therefore ev is a local isomorphism. \square

4.59 Remark

By definition 4.51 and 4.55, the set of points in general position is equal to the set of points in ev -general position. Recall that the points in ev -general position are the complement of a polyhedral complex of codimension 1, that is, they form a finite number of top-dimensional regions separated by “walls” that are polyhedra of codimension 1.

Due to lemma 4.58 it is sufficient for the proof of theorem 4.53 to consider a general point on such a wall and show that $\mathcal{P} \mapsto N_{\text{trop}}^{\text{irr}}(\Delta, g, \mathcal{P})$ is locally constant at these points, too. Such a general point on a wall is simply the image under ev of a general plane tropical curve C of a combinatorial type α such that $\dim \mathcal{M}_{\text{trop}, g, n}^{\alpha}(\Delta) = 2n - 1$.

So we have to check that $N_{\text{trop}}^{\text{irr}}(\Delta, g, \mathcal{P})$ is locally constant around such a point $C \in \widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta)$. That is, we have to show that for a neighborhood $U(C)$ around the point $C \in \widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta)$, the sum of the multiplicities of the preimages C' in $U(C)$, $\sum_{C' \in \text{ev}^{-1}(\mathcal{P}'') \cap U(C)} \text{mult } C'$, stays constant and does not depend on the point \mathcal{P}'' . (We will denote this sum in the following still by “ $N_{\text{trop}}^{\text{irr}}(\Delta, g, \mathcal{P})$ locally around C ”, even if we do not count the whole sum here, only the summands close to C .)

By proposition 4.41, we know that the types α such that $\dim \mathcal{M}_{\text{trop}, g, n}^{\alpha}(\Delta) = 2n - 1$ are precisely the types of codimension 1 and the exceptional types.

So we have to check that locally at a point $C \in \mathcal{M}_{\text{trop}, g, n}^{\alpha}(\Delta)$, where α is a type of codimension 1 or an exceptional type, $N_{\text{trop}}^{\text{irr}}(\Delta, g, \mathcal{P})$ stays constant. There are three cases we have to check:

- (1) C is of genus g and has exactly one 4-valent vertex,
- (2) C is of genus $g' = g - 1$ and is 3-valent or
- (3) C is of an exceptional type.

For the last two cases, we can immediately see that the sum over all inverse images of ev stays constant:

4.60 Lemma

Let C be a 3-valent (relevant) curve of genus $g - 1$ (that is, C is as in case 2 of remark 4.59).

Then $\mathcal{P} \mapsto N_{\text{trop}}^{\text{irr}}(\Delta, g, \mathcal{P})$ is locally constant around C .

Proof:

Let $\mathcal{P}' = \text{ev}(C)$. We have to see that for a configuration \mathcal{P}'' in general position near \mathcal{P}' , the sum of the inverse images under ev near C (counted with multiplicity) is constant and does not depend on \mathcal{P}'' . Let \mathcal{P}'' (near \mathcal{P}') be in general position. Then the inverse images under ev of \mathcal{P}'' near C are all 3-valent curves of genus g , such that C is in the boundary of their parametrizing spaces $\mathcal{M}_{\text{trop}, g, n}^{\alpha'}$. Let C' be such an inverse image. C' is of genus g and has therefore a loop which must disappear to a vertex of C . As C' is 3-valent and relevant, it cannot be a flat loop. But then the three edges adjacent to the vertex of C to which the loop disappears have to span \mathbb{R}^2 . In particular, none can be a marked point. That is, in C' there cannot be a marked point adjacent to the vanishing loop. Therefore C' is not rigid, and counts with multiplicity 0. Hence no matter which types we have in the inverse image of \mathcal{P}'' , the sum of the multiplicities is always 0. \square


4.61 Lemma

Let C be a (relevant) curve of an exceptional type α (that is, C is as in case 3 of remark 4.59).

Then $\mathcal{P} \mapsto N_{\text{trop}}^{\text{irr}}(\Delta, g, \mathcal{P})$ is locally constant around C .


Proof:

Let $\mathcal{P}' = \text{ev}(C)$. First recall that by 4.31 C is in the boundary of all types where the flat loop is “resolved”, that is, where the graph contains the following graph as subgraph:



So let C' be a curve of such a type which contains C in its boundary. We can assume that C' is rigid, as otherwise it would count 0. Then there must be at least one marked point adjacent to the bounded edges e and e' , because the loop must be broken. It is also possible that there are two marked points, one adjacent to e , the other to e' . (There cannot be more marked points, as otherwise C' would not be rigid again. The remaining part of C' would contain a string as it is not marked by enough points, see 4.50.)

Let us first assume there is one marked point x_i adjacent to the flat loop. Then we know that there are two types α_1 and α_2 that have α in their boundary. They are shown in the following picture.



The multiplicities of a curve of type α_1 obviously coincides with the multiplicity of a curve of type α_2 . So we only have to see that for a configuration \mathcal{P}'' in general position near \mathcal{P}' , a curve of exactly one of the types appears as preimage. As C is in the boundary of a rigid type, the unbounded edges we can reach via e_2, \dots, e_5 have to be separated by the marked points. We conclude that from x_i , we can reach an unbounded edge (without meeting marked points) by a path in Γ which involves precisely one of the four edges e_2, \dots, e_5 adjacent to the flat loop. Without restriction, let us assume that this is e_2 . That is, the lines to which the other three edges e_3, e_4 and e_5 are mapped to are fixed by the marked points $\mathcal{P}' \setminus \{p'_i\}$. This holds of course too for a curve of type α_1 and α_2 and the marked points $\mathcal{P}'' \setminus \{p''_i\}$. So the question which type appears depends on whether the point p''_i of the configuration \mathcal{P}'' lies above or below the line through $h(e_4) \cap h(e_5)$ with direction v_1 . (Due to the balancing condition, the direction vectors v_4 and v_5 of e_4 and e_5 must point to different sides of the line generated by v_1 , as indicated in the picture.)

If there are two marked points, we can use similar arguments. Two of the edges will be fixed by other marked points, and the question which type appears depends on which of the two points lies above the other (with respect to a line of direction v_1). \square

It would in fact be possible to give a similar (even though more complicated) argument for the first case of remark 4.59 (see theorem 4.8 of [14]). Here, we prefer to give another proof, which takes advantage of the structure of a polyhedral complex.

We define locally around C a space obtained by gluing some polyhedra, and a morphism of this space. The latter imitates the evaluation map but has the advantage that the absolute value of its determinant (in the maximal cells) is equal to the multiplicity of a corresponding tropical curve.

The new space that we are about to define, as well as the morphism, is defined by gluing pieces corresponding to the type of C and the types which have C in their boundary. We will start by defining these pieces and the maps on them and then define how to glue them.

4.62 Definition

Let α be a combinatorial type. Choose coordinates of $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ — that is, a root vertex V and an order of the bounded edges. Define the polyhedron $\widehat{\mathcal{M}}_{\text{trop}, g, n}^\alpha(\Delta)$ to be the subset of $\mathbb{R}^{2+\#\Gamma_0^1}$ given by the linear strict inequalities that all coordinates that correspond to lengths of edges have to be positive.

A different choice of the coordinates of $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ only leads to an isomorphism of $\widehat{\mathcal{M}}_{\text{trop}, g, n}^\alpha(\Delta)$ of determinant ± 1 as we have seen in 4.21.

Recall that a type of genus g and degree Δ has $n + \#\Delta - 3 + 3g - \sum_V(\text{val } V - 3) = 2n - 2 + 2g - \sum_V(\text{val } V - 3)$ bounded edges. (Remark 4.25 shows this for a 3-valent graph. Each higher valent vertex can be “resolved” by adding $\text{val } V - 3$ edges.) In particular, for a type of codimension 0, we have $\dim \widehat{\mathcal{M}}_{\text{trop}, g, n}^\alpha(\Delta) = 2n + 2g$.

Note that $\widehat{\mathcal{M}}_{\text{trop}, g, n}^\alpha(\Delta)$ contains the space $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ as the subset given by the conditions that the loops close up.

4.63 Definition

For a type α we choose the following data:

- for each marked point x_i a chain of flags leading from the root vertex V to x_i and
- a set of generators of $H_1(\Gamma, \mathbb{Z})$, where each such generator is given as a chain of flags around the loop.

Depending on these choices, we define a linear map $f_\alpha : \widehat{\mathcal{M}}_{\text{trop}, g, n}^\alpha(\Delta) \rightarrow \mathbb{R}^{2 + \#\Gamma_0^1}$ by defining a $2n + 2g'$ times $2 + \#\Gamma_0^1$ matrix with

- for each marked point two rows with the linear equation describing the position of $h(x_i)$ (depending on the position of $h(V)$ and the lengths of the bounded edges in the chosen chain of flags from V to x_i):

$$h(V) + \sum_F v(F) \cdot l([F]),$$

where the summation goes over all flags F in the chosen chain from V to x_i ; and

- for each chosen generator of $H_1(\Gamma, \mathbb{Z})$ two rows with the equation of the loop (depending on the lengths of the bounded edges in the loop):


$$\sum_F v(F) \cdot l([F]),$$

where F now goes over a chain of flags around the loop.

Note that f_α is a square matrix if and only if the number of marked points is equal to the number of nonmarked unbounded edges plus the genus minus one.

4.64 Example

The picture shows a tropical curve of genus 1.



The chosen root vertex is denoted by V , the order of the bounded edges is indicated, and the directions of the edges which are prescribed by the type are labelled. For the marked point x_1 , we choose to go from V via the bounded edges l_6 and l_1 , of direction v_6 respectively v_1 . For x_2 , we choose to go from V via the bounded edges l_6 , l_2 and l_3 , of direction v_6 , v_2 and v_3 . For x_3 , we choose to go via l_5 of direction $-v_4$. There is no choice for the generator of $H_1(\Gamma, \mathbb{Z})$. The equation given by a chain of flags around the unique loop is $l_6 \cdot v_6 + l_2 \cdot v_2 + l_4 \cdot v_4 + l_5 \cdot v_4 = 0$.

Hence the matrix of size 8 by 8 of f_α as defined above is:

$$\begin{pmatrix} E_2 & v_1 & 0 & 0 & 0 & 0 & v_6 \\ E_2 & 0 & v_2 & v_3 & 0 & 0 & v_6 \\ E_2 & 0 & 0 & 0 & 0 & -v_4 & 0 \\ 0 & 0 & v_2 & 0 & v_4 & v_4 & v_6 \end{pmatrix}$$

where each row stands for two rows, either corresponding to a marked point or to a loop.

4.65 Remark

Note that the map f_α of definition 4.63 depends on the choices we made, while the absolute value of the determinant of f_α does not. First, a coordinate change for $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ has determinant ± 1 and will therefore leave the absolute value of $\det f_\alpha$ unchanged. The same holds of course for a coordinate change of the target space $\mathbb{R}^{2n+2g'}$, that is, a different order of the marked points and loops.

If there are two chains of flags from a root vertex V to a marked point x_i , then their difference is a loop. (In the example above, we could also go from V to x_2 over l_5 , l_4 and l_3 with directions $-v_4$, $-v_4$ and v_3 .) Assume first that this loop is one of our chosen generators of $H_1(\Gamma, \mathbb{Z})$. Then choosing one or the other chain of flags from above just corresponds to adding (respectively subtracting) the two equations of the loop from the two rows of the marked point. (In the example, we have to subtract the two last rows if we choose the different chain of flag to x_2 .)

If we choose another set of generators for $H_1(\Gamma, \mathbb{Z})$, these new generators are given as linear combinations with coefficients in \mathbb{Z} of the old generators. Therefore, we also get linear combinations with coefficients in \mathbb{Z} of the equations. Hence a different choice of the generators of $H_1(\Gamma, \mathbb{Z})$ does either not change the determinant.

Therefore, we also get no different result if we choose for a marked point x_i another chain of flags, such that the difference is a loop that is not one of our chosen generators.

By abuse of notation, we will still speak of the map f_α , even though its definition depends on the choices we made, and keep in mind that $|\det(f_\alpha)|$ is uniquely determined, no matter what choices we made.

4.66 Remark

In definition 4.62 we enlarged the spaces $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ by ignoring the conditions given by the loops. But the map f_α from 4.63 still carries the information about the conditions:

the inverse image

$$f_\alpha^{-1}(\mathbb{R}^{2n} \times \{0\}) \subset \widehat{\mathcal{M}}_{\text{trop}, g, n}^\alpha(\Delta)$$

of those points in $\mathbb{R}^{2n+2g'}$ which have zeros at the entries to which the equations of the loops map to is equal to the space $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$.

Note that for a point $(\mathcal{P}, 0) \in \mathbb{R}^{2n} \times \{0\}$ a preimage (under f_α) $C \in \mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ is a tropical curve such that the markings are sent to the points \mathcal{P} , that is $\text{ev}(C) = \mathcal{P}$. That is, the map f_α from the larger space where we ignored the conditions of the loops can also be thought of as an enlarged evaluation map that carries the information about the loops.

What is important about the maps f_α is that for relevant 3-valent types, its determinant coincides with the multiplicity of a curve C of type α . This will be shown in the following two lemmata.

4.67 Lemma

Let C be a relevant 3-valent curve of genus g , degree Δ and type α , which is marked by $\#\Delta + g - 1$ points. Then the determinant $\det f_\alpha$ (where f_α is defined in 4.63) is zero if and only if $\text{mult } C = 0$ (where $\text{mult } C$ is defined in 4.47).

Proof:

By proposition 4.49 we know that the multiplicity of C is zero if and only if the evaluation map $\text{ev}|_{\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)}$ is not injective. But $\text{ev}|_{\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)}$ has more than one preimage for a point $\mathcal{P} \in \mathbb{R}^{2n}$ if and only if also $f_\alpha^{-1}(\mathcal{P}, 0)$ consists of more than one element (see remark 4.66). Hence $\det f_\alpha = 0$ if and only if $\text{mult } C = 0$. \square

4.68 Lemma


Let C be a relevant 3-valent curve of genus g , degree Δ and type α , which is marked by $\#\Delta + g - 1$ points. Then $|\det f_\alpha|$ (where f_α is defined in 4.63) is equal to $\text{mult } C$ (where $\text{mult } C$ is defined in 4.47).

Proof:

By lemma 4.67 we know that the statement is true if $\text{mult } C = 0$. Therefore we can now assume $\text{mult } C \neq 0$, in particular, we can assume that C is rigid.

The proof will be an induction on the sum of the number of bounded edges and the genus.

As induction beginning, we need a curve which is 3-valent, has a vertex to which no marked point is adjacent and where the sum of the number of bounded edges and the genus is minimal. As we need a vertex with no marked points, we need at least 3 unbounded edges which are not mapped to a point. As C is assumed to be rigid, it must have at least two marked points then. Therefore the following curve is adequate for the induction beginning:



The direction vectors of the two bounded edges are denoted by u and v . The curve is rational, therefore the matrix we have to set up to compute $\det f_\alpha$ contains only the rows corresponding to the evaluation map. We choose the root vertex to be V . Then the marked point x_1 is mapped to $h(V) + l_1 \cdot u$, where l_1 denotes the length of the bounded edge leading to x_1 . The marked point x_2 is mapped to $h(V) + l_2 \cdot v$. Therefore the matrix describing f_α is

$$A := \begin{pmatrix} E_2 & u & 0 \\ E_2 & 0 & v \end{pmatrix}$$

A computation shows $|\det A| = |\det(u, v)|$. The latter is by definition the multiplicity of C .


Note that there is no other rigid curve where the sum of the number of bounded edges and the genus is 2 or less. So this computation is enough for the induction beginning.

As induction step, let us now assume C has k bounded edges, is a curve of genus g and degree Δ , and $k + g > 2$. Cut one of the bounded edges. That is, in the graph Γ , choose a bounded edge e and replace it by two unbounded edges, each being adjacent to one end vertex of e . Two things can happen:

- (1) The graph could decompose into two connected components that we will denote by Γ_1 and Γ_2 . In this case, the edge e should be chosen such that both Γ_1 and Γ_2 contain at least one bounded edge.
- (2) The graph could stay connected, but a loop could be broken. We denote the new connected graph of genus $g - 1$ by Γ_1 . In this case, the edge e should be chosen such that it is adjacent to a marked point. (Such a choice is possible as C is rigid.)

We have to prove the statement for each of the two cases separately, as the arguments differ:

For the first case, denote the two tropical curves that arise when we replace e by two unbounded edges with C_1 and C_2 . Assume C_1 is of type α_1 and C_2 of type α_2 .



Let us assume C_1 has $e_1 + 1$ unbounded edges which are not marked points, and C_2 has $e_2 + 1$ unbounded edges (that is, $e_1 + e_2 = \#\Delta$ is the number of unbounded edges of C , and the cut edge e counts as a new unbounded edge both for C_1 and C_2). Let us also assume that the genus of C_1 is equal to g_1 and the genus of C_2 is g_2 . As the cutting of

e leads to a nonconnected graph, it cannot have broken a loop. Therefore $g_1 + g_2 = g$ is the genus of C . Let us now assume C_1 has $l_1 < (e_1 + 1) + g_1 - 2$ marked points. That is, we assume that there are at least two points less than we need to separate all unbounded edges and loops. Then by lemma 4.50 C_1 would not be rigid. If there was a loop which does not contain a marked point, the same would hold for C , and so C would not be rigid. If there were two ends which are connected by a string, then one of these ends could of course be the new unbounded edge which replaces e . So we cannot immediately conclude that C contains a string, too. Let us mark this new unbounded edge of C_1 , too. Then we have $l_1 + 1 < (e_1 + 1) + g_1 - 1$ marked points, which is still not enough to separate the ends and loops. Hence there is still a string which is contained in C , too.

So we have $l_1 \geq e_1 + g_1 - 1$, and analogously, also $l_2 \geq e_2 + g_2 - 1$. But $l_1 + l_2$ is the number of marked points of C , which is equal to $n = \#\Delta + g - 1 = e_1 + g_1 + e_2 + g_2 - 1$. Therefore there are only two possibilities: $l_1 = e_1 + g_1$ and $l_2 = e_2 + g_2 - 1$ or vice versa. Without loss of generality, let $l_1 = e_1 + g_1$ and $l_2 = e_2 + g_2 - 1$.

Let us now describe the matrix f_α . We choose the root vertex to be the vertex in ∂e that belongs to C_1 after cutting. Then we begin with the marked points and loops which belong to C_1 . These marked points and loops need only the coordinates of the root vertex (which is chosen to lie in C_1) plus the coordinates of C_1 . As C_1 as well as C is 3-valent, we have by remark 4.25 $e_1 + l_1 - 2 + 3g_1 = 2e_1 + 4g_1 - 2$ bounded edges in C_1 .

So for the $l_1 = e_1 + g_1$ marked points and g_1 loops of C_1 , we need $2 + 2e_1 + 4g_1 - 2 = 2e_1 + 4g_1$ coordinates. That is, the matrix is of the form

$$\left(\begin{array}{c|c} A_1 & 0 \\ \hline * & A_2 \end{array} \right)$$

where A_1 and A_2 are square matrices of size $2e_1 + 4g_1$ and $2e_2 + 4g_2 - 2$. A_1 is actually the matrix f_{α_1} . C_1 has fewer bounded edges than C , and we can therefore by induction assume that $\text{mult } C_1 = |\det f_{\alpha_1}| = |\det A_1|$. (We choose for C_1 the same root vertex, order of bounded edges, chains of flags from the root vertex to the marked points and loops.) So $\det f_\alpha = \det f_{\alpha_1} \cdot \det A_2$. C_2 is now marked with $l_2 = e_2 + g_2 - 1$ points, and is therefore not rigid. More precisely, we know that e must be part of the string of C_2 , because C is rigid. Therefore we mark e by another marked point, and choose the root vertex of C_2 to be the new vertex adjacent to this marked point. The curve C_2 together with this new marked point has again fewer bounded edges than C , and we can by induction assume that $\text{mult } C_2 = |\det f_{\alpha'_2}|$ (where α'_2 denotes now the type of C_2 together with the new marked point).

The matrix A_2 contains the rows that describe the position of the marked points of C_2 depending on the root vertex in C_1 . Every chain of flags from the root vertex in C_1 that leads to a marked point in C_2 has to pass e . Furthermore, A_2 contains the rows that describe equations of loops in C_2 .

Let us now compare A_2 with the matrix $f_{\alpha'_2}$ of C_2 with the newly added marked point. The equations of the loops in C_2 that occurred in A_2 will also occur in the matrix $f_{\alpha'_2}$ — only we have to add zeros for the coordinates of the new root vertex. The rows that

describe the position of the marked points in C_2 now depend on the new root vertex. But as this is chosen to be adjacent to the marked point on e , every chain of flags from this root vertex to a marked point has to pass e , too. So we can choose the same chains of flags as for C . As we chose the root vertex to be the vertex adjacent to the new marked point, the position of this marked point coincides with the position of the root vertex. That is, the matrix $f_{\alpha'_2}$ is of the following form:

$$\begin{pmatrix} E_2 & 0 \\ E_2 & \\ \vdots & \\ E_2 & A_2 \\ 0 & \\ \vdots & \\ 0 & \end{pmatrix}$$


Therefore, $\det A_2$ is indeed equal to the determinant of the matrix $f_{\alpha'_2}$.

To sum up, we have seen that

$$|\det f_\alpha| = |\det f_{\alpha_1}| \cdot |\det f_{\alpha'_2}| = \text{mult } C_1 \cdot \text{mult } C_2 = \text{mult } C,$$

where the second equality holds by induction and the last equality holds by definition of the multiplicity.

For the second case, let Γ_1 be the — still connected — graph that arises after cutting the edge e . The graph Γ_1 has genus $g-1$, it has $\#\Delta + 2$ unbounded edges that are not marked points, and it has $\#\Delta + g - 1 < (\#\Delta + 2) + (g - 1) - 1$ marked points, therefore it cannot be rigid by 4.50. We add a marked point x adjacent to one of the new unbounded edges. There is only one possibility to choose this unbounded edge such that the new tropical curve is rigid. Recall that we chose e such that it is adjacent to a marked point x_i .



The tropical curve C_1 of type α_1 defined in this way has genus $g-1$ and as many bounded edges as C . Therefore we can assume by induction that its multiplicity is equal to $|\det f_{\alpha_1}|$. As the multiplicity of C is equal to the multiplicity of C_1 , it remains to see that $|\det f_\alpha| = |\det f_{\alpha_1}|$. So let us choose coordinates to compare the two matrices of f_α and f_{α_1} . Choose V — the vertex adjacent to the marked point x_i — as root vertex both for C and for C_1 . Choose the same order of bounded edges, marked points and loops for the two curves. One of the loops, say L , of C is broken after the cutting of e . This loop corresponds to the

last two lines of the matrix of f_α . For C_1 , the last two lines shall be given by the marked point x . As chain of flags leading from V to x in C_1 , we choose just the same chain of flags as for the loop L . The following table represents both matrices. The two matrices only differ by the $h(V)$ -entries in the last two rows. In the table, each row represents two or more rows as before. Each matrix contains the first three rows, f_α contains the fourth, and f_{α_1} the fifth.

	$h(V)$	bounded edges
the marked point x_i	E_2	0
other marked points	E_2	*
other loops	0	*
for f_α the loop L	0	equation for L
for f_{α_1} the new point x	E_2	equation for L

Note that both matrices are block matrices with a 2×2 block on the top left. Therefore, both determinants are equal to the determinant of the lower right block. But this block coincides for both matrices, because it does not involve the two numbers we changed from 0 to 1.

Hence $|\det f_\alpha|$ is equal to $|\det f_{\alpha_1}|$, and the latter is seen by induction to be equal to $\text{mult } C_1 = \text{mult } C$.

□

It would be desirable to define the polyhedral complex $\widehat{\mathcal{M}}_{\text{trop}, g, n}(\Delta)$ with cells $\widehat{\mathcal{M}}_{\text{trop}, g, n}^\alpha(\Delta)$ and the morphism of polyhedral complexes $f : \widehat{\mathcal{M}}_{\text{trop}, g, n}(\Delta) \rightarrow \mathbb{R}^{2n+2g}$ globally. Then for theorem 4.53 it would be enough to prove that the map $\mathcal{P} \mapsto \deg_f((\mathcal{P}, 0))$ is constant (even without considering the two other special cases in lemma 4.60 and lemma 4.61). But we have given the map only locally as f_α on the sets $\widehat{\mathcal{M}}_{\text{trop}, g, n}^\alpha(\Delta)$, and even though $|\det f_\alpha|$ is unique, the map f_α does depend on the choices (see remark 4.66). There is no “global” description of this map as for the evaluation map, because the elements of $\widehat{\mathcal{M}}_{\text{trop}, g, n}^\alpha(\Delta)$ are not even tropical curves, only those are where the conditions that the loops close up are fulfilled. So for different types, which contain the same type in their boundary, we would have to make a “consistent” choice of the coordinates in order to get maps f_α that coincide on the boundary. Due to the large amount of different types and their boundary relations, this aim seems unachievable.

We will therefore only define a local space $\widehat{\mathcal{M}}$ over a general point on a codimension 1-“wall” as in remark 4.59, case 1, and locally a map f from this space. Our aim is then to show that the degree of f is locally constant.


4.69 Remark

Note that for the case of rational curves, the spaces $\widehat{\mathcal{M}}_{\text{trop}, g, n}^\alpha(\Delta)$ coincide with the spaces $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$, and f_α is just the evaluation map. In this case, we therefore have a global space and a global description of f , and we can prove theorem 4.53 (without using lemma

4.60 and lemma 4.61, as these cases cannot appear for rational curves) by showing that the map $\mathcal{P} \mapsto \deg_f(\mathcal{P})$ is constant. (Note that lemma 4.68 shows that for rational curves $\text{mult}_{\text{ev}}(C)$ is equal to $\text{mult}(C)$.) For this, it will as before be sufficient to prove that $\mathcal{P} \mapsto \deg_f(\mathcal{P})$ is locally constant which follows with the same “local” proof as below for lemma 4.72.

4.70 Definition

Let C be a tropical curve of a type α of genus g and with exactly one 4-valent vertex, as in case 1 of remark 4.59. From remark 4.31 we know that there are three types which contain a curve of type α in the boundary. We will denote them as in the picture with α_1, α_2 and α_3 . They differ in which of the four edges come together at a 3-valent vertex, and in the direction of the new edge e that separates the 4-valent vertex.



For all i , we take the polyhedra $\widehat{\mathcal{M}}_{\text{trop}, g, n}^{\alpha_i}(\Delta)$, and also $\widehat{\mathcal{M}}_{\text{trop}, g, n}^{\alpha}(\Delta)$.

Choose the root vertex to be V as in the picture. Choose an order of the bounded edges which coincides for all types at the edges that appear in all types, and such that the edge e comes last. For all types and all marked points, choose a chain of flags from V to the marked point that may differ only at e . Also, choose a set of generators of the loops for each type which may differ only at e .

We identify the part of the boundaries of $\widehat{\mathcal{M}}_{\text{trop}, g, n}^{\alpha_i}(\Delta)$ that corresponds to curves where the coordinate of e is 0 with $\widehat{\mathcal{M}}_{\text{trop}, g, n}^{\alpha}(\Delta)$. The space defined in that way is denoted by $\widehat{\mathcal{M}}$.

For the coordinates as chosen above, the four maps $f_{\alpha}, f_{\alpha_1}, f_{\alpha_2}$ and f_{α_3} coincide at the identified points and we can therefore glue them to one map $f : \widehat{\mathcal{M}} \rightarrow \mathbb{R}^{2n+2g}$.

4.71 Remark

The space $\widehat{\mathcal{M}}$ from definition 4.70 can be thought of as a “local polyhedral complex” of pure dimension $2n + 2g$, and the map

$$f : \widehat{\mathcal{M}} \rightarrow \mathbb{R}^{2n+2g}$$

as a “local morphism of polyhedral complexes of the same dimension”. $\widehat{\mathcal{M}}$ consists of four cells (however, their closures are not completely contained in $\widehat{\mathcal{M}}$). Three of the cells are of dimension $2n + 2g$, the one which is of lower dimension is contained in their boundary. The coordinate change is obviously a linear map, as it just drops one coordinate (corresponding to the length of the bounded edge e , as in 4.70). $f : \widehat{\mathcal{M}} \rightarrow \mathbb{R}^{2n+2g}$ maps each cell to the one open cell of \mathbb{R}^{2n+2g} , which is a polyhedral complex. By abuse of notation, we will therefore speak of the degree of f , even though it is not really a morphism of polyhedral complexes.

We are finally ready to prove the last statement which is missing for the proof of theorem 4.53:

4.72 Lemma

Let C be a (relevant) curve with exactly one 4-valent vertex, and of genus g (that is, C is as in case 3 of remark 4.59).

Then $\mathcal{P} \mapsto N_{\text{trop}}^{\text{irr}}(\Delta, g, \mathcal{P})$ is locally constant around C .

Proof:

Let $\text{ev}(C) = \mathcal{P}'$. Recall that $N_{\text{trop}}^{\text{irr}}(\Delta, g, \mathcal{P}) = \sum_{C' \in \text{ev}^{-1}(\mathcal{P})} \text{mult } C'$. By remark 4.66, we have

$$(\text{ev} |_{\mathcal{M}_{\text{trop}, g, n}^{\alpha}(\Delta)})^{-1}(\mathcal{P}) = f_{\alpha}^{-1}(\mathcal{P}, 0).$$

In particular, for the four types that we glued to get $\widehat{\mathcal{M}}$ in definition 4.70, and for a configuration \mathcal{P}'' near \mathcal{P}' we have:

$$\text{ev}^{-1}(\mathcal{P}'') \cap (\mathcal{M}_{\text{trop}, g, n}^{\alpha}(\Delta) \cup \bigcup_i \mathcal{M}_{\text{trop}, g, n}^{\alpha_i}(\Delta)) = f^{-1}(\mathcal{P}'', 0)$$

Therefore it is enough to see that

$$\mathcal{P} \mapsto \sum_{C' \in f^{-1}(\mathcal{P}, 0)} \text{mult } C'$$

is locally constant at \mathcal{P}' .

From lemma 4.68 we know that $\text{mult } C' = |\det f|$. Note that by remark 4.65 the latter does not depend on any choice of coordinates of $\widehat{\mathcal{M}}_{\text{trop}, g, n}^{\alpha_i}(\Delta)$, nor on the choice of the map f_{α_i} on each cell.


Hence,

$$\sum_{f^{-1}(\mathcal{P}, 0)} \text{mult } C' = \sum_{f^{-1}(\mathcal{P}, 0)} \text{mult}_f C' = \text{deg}_f(\mathcal{P}, 0)$$

by definition 4.55 and the latter is well-defined. (Recall that by 4.71 we use the notation of degree, even though $\widehat{\mathcal{M}}$ is not really a polyhedral complex.)

So our aim is now to show that $\mathcal{P} \mapsto \text{deg}_f(\mathcal{P}, 0)$ is locally constant at \mathcal{P}' .

To see this, we study the three matrices A_1 , A_2 and A_3 of f_{α_1} , f_{α_2} and f_{α_3} . They differ only in the column corresponding to the edge e . Denote the four edges adjacent to the 4-valent vertex of C with e_1 , e_2 , e_3 and e_4 , and their respective directions with v_1, \dots, v_4 . The root vertex is as before V as indicated in the picture.



As we assumed in definition 4.70 that all choices are made consistently, the three matrices only differ in the column which belongs to the new edge e . The following table represents all three matrices: Each matrix A_i contains the first block of columns (corresponding to

the image $h(V)$ of the root vertex and the lengths l_i of the edges e_i) and the i -th of the last three columns (corresponding to the length of the new edge e). The columns corresponding to the other bounded edges are not shown; it suffices to note here that they are the same for all three matrices. The first four rows correspond to the images of the marked points, the last six rows correspond to the equations of the loops. We get four different types of rows for the marked points depending on via which of the four edges e_i a marked point is reached from V . (Each row represents in fact two or more rows of the matrix, two rows for the two coordinates of the image of each marked point). For the loops, we get six different types of rows depending on which two of the four edges e_1, \dots, e_4 are involved in a loop. (Again, each row represents two or more rows for the two equations given by each loop.) Loops that do not involve any of the four edges are not added, they just have zeros in every of the special columns we show, and they do not change the computations.

	$h(V)$	l_1	l_2	l_3	l_4	l^{α_1}	l^{α_2}	l^{α_3}
points behind e_1	E_2	v_1	0	0	0	0	0	0
points behind e_2	E_2	0	v_2	0	0	0	$v_2 + v_3$	$v_2 + v_4$
points behind e_3	E_2	0	0	v_3	0	$v_4 + v_3$	$v_2 + v_3$	0
points behind e_4	E_2	0	0	0	v_4	$v_4 + v_3$	0	$v_2 + v_4$
loops involving e_1 and e_2	0	$-v_1$	v_2	0	0	0	$v_2 + v_3$	$v_2 + v_4$
loops involving e_1 and e_3	0	$-v_1$	0	v_3	0	$v_3 + v_4$	$v_2 + v_3$	0
loops involving e_1 and e_4	0	$-v_1$	0	0	v_4	$v_3 + v_4$	0	$v_2 + v_4$
loops involving e_2 and e_3	0	0	$-v_2$	v_3	0	$v_3 + v_4$	0	$-v_2 - v_4$
loops involving e_2 and e_4	0	0	$-v_2$	0	v_4	$v_3 + v_4$	$-v_2 - v_3$	0
loops involving e_3 and e_4	0	0	0	$-v_3$	v_4	0	$-v_2 - v_3$	$v_2 + v_4$

As det is linear in each column, we have $\det A_1 + \det A_2 + \det A_3$ is equal to the determinant of the following matrix, where we added the three last columns:

	$h(V)$	l_1	l_2	l_3	l_4	
points behind e_1	E_2	v_1	0	0	0	0
points behind e_2	E_2	0	v_2	0	0	$2v_2 + v_3 + v_4$
points behind e_3	E_2	0	0	v_3	0	$2v_3 + v_2 + v_4$
points behind e_4	E_2	0	0	0	v_4	$2v_4 + v_3 + v_2$
loops involving e_1 and e_2	0	$-v_1$	v_2	0	0	$2v_2 + v_3 + v_4$
loops involving e_1 and e_3	0	$-v_1$	0	v_3	0	$2v_3 + v_2 + v_4$
loops involving e_1 and e_4	0	$-v_1$	0	0	v_4	$2v_4 + v_2 + v_3$
loops involving e_2 and e_3	0	0	$-v_2$	v_3	0	$v_3 - v_2$
loops involving e_2 and e_4	0	0	$-v_2$	0	v_4	$v_4 - v_2$
loops involving e_3 and e_4	0	0	0	$-v_3$	v_4	$v_4 - v_3$

Now we subtract the four columns for l_1, \dots, l_4 from the last column.

	$h(V)$	l_1	l_2	l_3	l_4	
points behind e_1	E_2	v_1	0	0	0	$-v_1$
points behind e_2	E_2	0	v_2	0	0	$v_2 + v_3 + v_4$
points behind e_3	E_2	0	0	v_3	0	$v_3 + v_2 + v_4$
points behind e_4	E_2	0	0	0	v_4	$v_4 + v_3 + v_2$
loops involving e_1 and e_2	0	$-v_1$	v_2	0	0	$v_2 + v_3 + v_4 + v_1$
loops involving e_1 and e_3	0	$-v_1$	0	v_3	0	$v_3 + v_2 + v_4 + v_1$
loops involving e_1 and e_4	0	$-v_1$	0	0	v_4	$v_4 + v_2 + v_3 + v_1$
loops involving e_2 and e_3	0	0	$-v_2$	v_3	0	0
loops involving e_2 and e_4	0	0	$-v_2$	0	v_4	0
loops involving e_3 and e_4	0	0	0	$-v_3$	v_4	0

Due to the balancing condition $v_1 + v_2 + v_3 + v_4 = 0$, and finally we can add v_1 times the $h(V)$ -columns to the last column and get a matrix with a zero column whose determinant is 0. Therefore $\det A_1 + \det A_2 + \det A_3 = 0$.

Note that we assume here that the edges e_i are in fact all bounded. If this is not true, the argument needs to be changed slightly. If e_i is unbounded, then there can be no marked points that can be reached from V via e_i . That is, we do not have the corresponding rows.

For a given $i \in \{1, 2, 3\}$ let us now determine whether the combinatorial type α_i occurs in the inverse image of a fixed point $(\mathcal{P}'', 0) \in \mathbb{R}^{2n+2g}$ near $(\mathcal{P}', 0)$. We may assume without loss of generality that the multiplicity of α_i is non-zero since other types are irrelevant for the statement of the proposition. So the restriction f_{α_i} of f to $\widehat{\mathcal{M}}_{\text{trop}, g, n}^{\alpha_i}(\Delta)$ is given by the invertible matrix A_i . There is therefore at most one inverse image point in $f_{\alpha_i}^{-1}((\mathcal{P}'', 0))$, which would have to be the point with coordinates $A_i^{-1} \cdot (\mathcal{P}'', 0)$. In fact, this point exists in $\widehat{\mathcal{M}}_{\text{trop}, g, n}^{\alpha_i}(\Delta)$ if and only if all coordinates of $A_i^{-1} \cdot (\mathcal{P}'', 0)$ corresponding to lengths of bounded edges are positive. By continuity this is obvious for all edges except the newly added edge e , because in the boundary curve $C = f_{\alpha_i}^{-1}((\mathcal{P}', 0))$ all these edges had positive length. We conclude that there is a point in $f_{\alpha_i}^{-1}((\mathcal{P}'', 0))$ if and only if the last coordinate (corresponding to the length of the newly added edge e) of $A_i^{-1} \cdot (\mathcal{P}'', 0)$ is positive. By Cramer's rule this last coordinate is $\det \widetilde{A}_i / \det A_i$, where \widetilde{A}_i denotes the matrix A_i with the last column replaced by $(\mathcal{P}'', 0)$. But note that \widetilde{A}_i does not depend on i since the last column was the only one where the matrices A_i differ. Hence whether there is a point in $f_{\alpha_i}^{-1}((\mathcal{P}'', 0))$ or not depends only on the sign of $\det A_i$: either there are such inverse image points for exactly those i where $\det A_i$ is positive, or exactly for those i where $\det A_i$ is negative. But by the above the sum of the absolute values of the determinants satisfying this condition is the same in both cases. Hence $\mathcal{P} \mapsto \deg_f(\mathcal{P}, 0)$ is locally constant at \mathcal{P}' . \square

Let us finally sum up again the arguments we collected for the proof of theorem 4.53:

Proof of theorem 4.53:

By lemma 4.58 we know that $\mathcal{P} \mapsto N_{\text{trop}}^{\text{irr}}(\Delta, g, \mathcal{P})$ is locally constant on the subset of points in general position. So we only have to check that this map is also locally constant at a general configuration of points which is not in general position. Remark 4.59 tells us which points these are: there are three cases to check. The first is checked in 4.72, the second in 4.60 and the third in 4.61. \square

4.73 Remark

Recall that we restricted to connected graphs Γ .

The numbers $N_{\text{trop}}(\Delta, g, \mathcal{P})$ can be defined analogously to 4.52, where we replace the space $\widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta)$ by the space $\widetilde{\mathcal{M}}'_{\text{trop}, g, n}(\Delta)$. The latter contains also tropical curves where the underlying graph Γ is not connected (see 4.17). All concepts we had for $\widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta)$ (evaluation maps and so on) can be defined for $\widetilde{\mathcal{M}}'_{\text{trop}, g, n}(\Delta)$, too. The analogous statement for $N_{\text{trop}}(\Delta, g, \mathcal{P})$ as in theorem 4.53 holds, too. To prove it, we have to study the dimensions of the strata $\mathcal{M}'_{\text{trop}, g, n}(\Delta)$ of the moduli space of not necessarily connected relevant parametrized tropical curves. When we do this, we will see that the strata of dimension $2n - 1$ are precisely of the same types as in the proof above.

Therefore, we can also talk about the numbers $N_{\text{trop}}(\Delta, g)$, knowing that they do not depend on the position of the points.

5. THE CORRESPONDENCE OF TROPICAL CURVES AND LATTICE PATHS

Chapter 2.3 suggests an (although not 1 : 1-) correspondence of tropical curves and Newton subdivisions of the Newton polygon. So far, we enumerated parametrized tropical curves passing through certain marked points. The duality of tropical curves and Newton subdivisions brings up the idea that it should be possible to enumerate Newton subdivisions instead. We start the chapter in section 5.1 with the description of basic properties of parametrized tropical curves in the dual setting of Newton subdivisions. As the dual Newton subdivision depends only on the image $h(\Gamma)$ of a parametrized tropical curve, we have to deal with these images, too. They (together with some information about weights) will be called *unparametrized tropical curves*. When passing from a parametrized tropical curve (Γ, h) to its image $h(\Gamma) \subset \mathbb{R}^2$, we may lose a lot of information about the curve. A possibility to avoid this is to work with *simple* parametrized tropical curves, a notion which will also be explained in section 5.1.

When we try to count Newton subdivisions instead of tropical curves, we have to keep in mind that we only have to enumerate those Newton subdivisions which are dual to a tropical curve that passes through a given set of points. In order to distinguish those, we are going to study the dual of marked tropical curves in section 5.2. Also, we have to make a restriction on the set of points, because we only want to deal with simple curves. Therefore we introduce the notion of *restricted general position* — only simple curves will pass through such a set of points. We will show that the subset of points in restricted general position in \mathbb{R}^{2n} is still of top dimension.

For a special point configuration in restricted general position, the marked tropical curves passing through this set will be dual to *lattice paths* in the Newton polygon. We will introduce this notion in section 5.3, and also define the multiplicity of a lattice path. We can then define the number of lattice paths counted with multiplicity as $N_{\text{path}}(\Delta, g)$ — these numbers will be the analogues of the numbers $N_{\text{trop}}(\Delta, g)$.

In section 5.4 we will finally prove that $N_{\text{path}}(\Delta, g) = N_{\text{trop}}(\Delta, g)$ (theorem 5.44) by defining a bijection between the set of lattice paths and the set of tropical curves through our chosen point configuration (both counted with multiplicity). It is important to note that the multiplicity of a lattice path is not equivalent to the multiplicity of a tropical curve. In fact, our bijection will assign several tropical curves to one lattice path in some cases. The multiplicity of the path corresponds to the sum of the multiplicities of the assigned tropical curves.

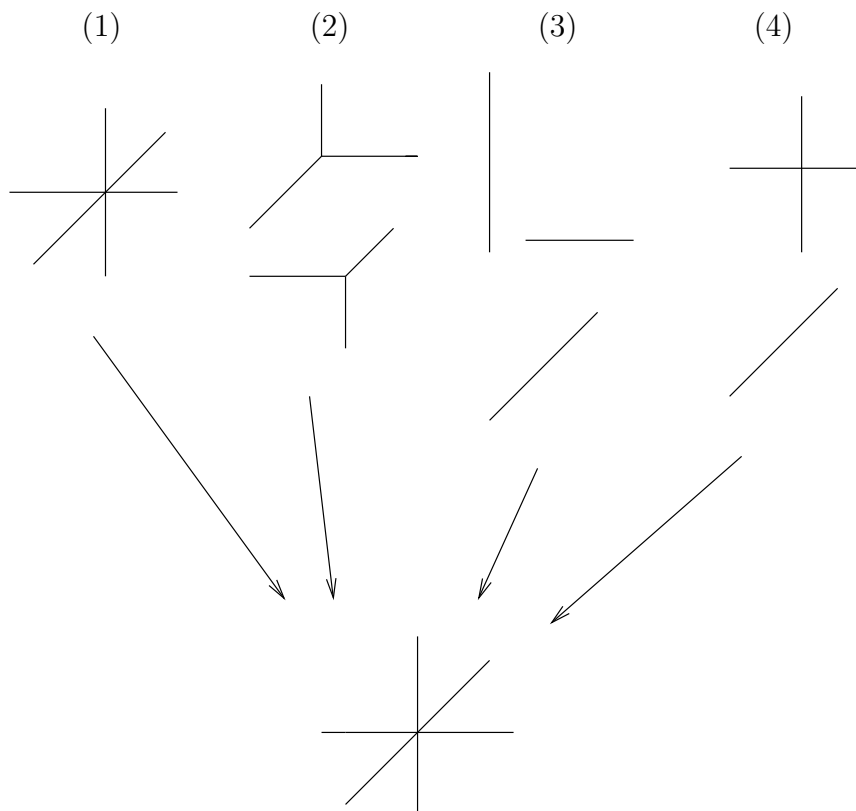
Contrary to the numbers $N_{\text{cplx}}(\Delta, g)$ and $N_{\text{trop}}(\Delta, g)$ the lattice paths numbers $N_{\text{path}}(\Delta, g)$ can be computed directly: we just have to take all λ -increasing paths in the polygon Δ and compute their multiplicity, which is defined recursively. In particular, the correspondence $N_{\text{path}}(\Delta, g) = N_{\text{trop}}(\Delta, g) = N_{\text{cplx}}(\Delta, g)$ allows us to compute all the numbers $N_{\text{cplx}}(\Delta, g)$ using the lattice path count.

The results explained in this chapter were achieved by Mikhalkin in [23]. However, we wish to present a proof here, too, adding some details and adjusting it to our definitions for parametrized curves, which slightly differ from the definitions in [23].

5.1. PROPERTIES OF PARAMETRIZED TROPICAL CURVES DESCRIBED IN THE DUAL LANGUAGE OF NEWTON SUBDIVISIONS

First recall how the dual of a tropical curve is constructed. Given a regular parametrized tropical curve (Γ, h) , we know that there is a curve over the completion of the field of Puiseux series K , such that $\overline{\text{Val}(C)} = h(\Gamma)$ by theorem 4.27. This curve is given by a polynomial f , and by theorem 2.9 we know that the tropical curve associated to it (that is, $h(\Gamma)$) is equal to the tropical curve associated to the tropical polynomial $\text{trop } f$. For the tropical curve associated to $\text{trop } f$ (again, that is for $h(\Gamma)$) we know by theorem 2.15 that it is dual to a subdivision of the Newton polygon of $\text{trop } f$.

Note that the dual Newton subdivision depends only on the image $h(\Gamma)$ of a parametrized tropical curve, not on the parametrization. Therefore we need to study these images first. Given an image $h(\Gamma) \subset \mathbb{R}^2$, there can be several possibilities to parametrize it. The following picture shows such an image and several possibilities how it could be parametrized. It could for example be parametrized by a graph with a 6-valent vertex (1), or by a disconnected graph with two 3-valent vertices (2). It is furthermore possible that the two vertices from (2) are connected by an edge which is mapped to a point. In (3) and (4), the connected components in the picture could also be linked by a contracted edge.



To avoid this ambiguity, we will first make an assumption on the parametrized curves we want to work with:

5.1 Definition

A parametrized tropical curve (Γ, h) is called *simple* if

- Γ is 3-valent,
- h is injective on the set of vertices Γ^0 of Γ ,
- the images $h(V)$ and $h(e)$ of a vertex $V \in \Gamma^0$ and an edge $e \in \Gamma^1$ are disjoint, that is $h(V) \cap h(e) = \emptyset$,
- the images $h(e_1)$ and $h(e_2)$ of two edges $e_1, e_2 \in \Gamma^1$, $e_1 \neq e_2$, have at most one point in common, that is $\#\{h(e_1) \cap h(e_2)\} \leq 1$ and
- through each point $p \in \mathbb{R}^2$ pass at most 2 edges, that is $\#h^{-1}(p) \leq 2$.

Note that each simple parametrized tropical curve is also relevant. There cannot be a contracted edge, because else this edge and its adjacent vertices would be mapped to the same image point. Also, there cannot be a vertex where two edges point in the same direction. (The images of these two edges would have more than one point in common.)

5.2 Remark


Let $h(\Gamma)$ be the image of a simple parametrized tropical curve. As we have already seen, no edge of Γ can be contracted to a point by h . Due to the third and fourth property of simple curves, no vertex can be mapped onto an edge and two edges cannot be mapped to the same line segment. Therefore we can distinguish the images of the vertices and edges in $h(\Gamma)$ — each edge of Γ is mapped to a line segment of the image (and no two edges on the same), and each vertex of Γ will be mapped to a point of the image where three line segments come together.

5.3 Definition

Let (Γ, h) be a simple parametrized tropical curve. As (Γ, h) is simple, we can distinguish the images of the edges in $h(\Gamma) \subset \mathbb{R}^2$ (see 5.2). Associate to each edge $h(e)$ of $h(\Gamma)$ its weight $\omega(e)$ (see 4.13). Then the image $h(\Gamma)$ together with the associated weights is called an *unparametrized tropical curve* $C = (h(\Gamma), \{(h(e), \omega(e))\})$.

5.4 Example

The following picture shows an unparametrized tropical curve. The weights equal to 1 are not marked in the picture. There are two edges of weight 2.



5.5 Remark

Let C be an unparametrized tropical curve. Then the simple parametrization (Γ, h) of C is uniquely determined (up to isomorphism). This is true because by 5.2 we can distinguish the images of the edges and vertices. Therefore the homeomorphism class of the graph Γ is uniquely determined by C . As we know the lines to which the edges are mapped, we can determine the primitive integral vectors $u(F)$ for each flag F . For each edge e , we also know the weight $\omega(e)$ (as the weights are marked for each edge in C), hence we can determine the directions $v(F)$ for each flag. The length $l(e)$ of an edge e of the graph is then determined by the length of its prescribed image $h(e)$ and the direction. Therefore,

not only the homeomorphism class, but also the graph Γ is uniquely given by C . The map h is then uniquely given, too, as it is linear on each edge.

5.6 Definition

We define the *genus* and *degree* of an unparametrized tropical curve to be the genus and degree of the unique simple parametrization (see 5.5).

Also, we define the *multiplicity* of an unparametrized tropical curve to be equal to the multiplicity of the simple parametrization.

An unparametrized tropical curve is called *irreducible*, if the graph Γ of the unique simple parametrization is connected, and *reducible* otherwise.

A *component* of a reducible unparametrized tropical curve is the image of a connected component of the graph Γ of the simple parametrization.


5.7 Remark

In all situations we will work with later on, the simple parametrization of an unparametrized tropical curve is unique also without the information about the weights of the edges.

Given a component of an unparametrized tropical curve, the information of one weight is enough to determine all other weights with the aid of the balancing condition: The unparametrized tropical curve is given as the image of a simple parametrized curve. For a simple parametrized tropical curve, we have only 3-valent vertices such that the three edges adjacent span \mathbb{R}^2 . If the weight of one edge is given, the two other weights can be determined uniquely.

In our definition of unparametrized tropical curve, we need to associate these weights, because otherwise there would be some ambiguity in the choice of a parametrization: we could take a multiple of the weights of all edges which belong to the same component.

In the picture below, we could parametrize the image by curve of degree 3 or of degree $\Delta := \{(-2, 0), 2 \cdot (-1, 0), (0, -2), 2 \cdot (0, -1), (2, 2), 2 \cdot (1, 1)\}$, for example.



Only if we fix the weights in the image as we have done it for unparametrized tropical curves, then there is no choice left.

However, we will later on only prescribe degrees such that at least all the unbounded edges which are mapped to lines of the primitive directions $(0, -1)$ and $(1, 1)$ are of weight 1. In each component of an unparametrized tropical curve of such a degree, there has to be an unbounded edge of weight 1. So at least one weight is prescribed in every component. This weight determines the other weights as we have seen above.

Hence, in the cases we will work with later, we can uniquely parametrize an unparametrized tropical curve, even without knowing the weights of the edges (that is, we can parametrize an image $h(\Gamma) \subset \mathbb{R}^2$ of a simple curve uniquely).

By abuse of notation, we will in the following sometimes denote an unparametrized tropical curve by $h(\Gamma)$, neglecting the information about the weights.

Let us now come back to the dual Newton subdivisions.

First, we want to make precise what Newton subdivision should mean.


5.8 Definition

Let Δ be a convex lattice polygon in \mathbb{R}^2 . Let $\Delta_1, \dots, \Delta_k$ be a collection of convex lattice polygons (given as convex hulls of their vertices in \mathbb{Z}^2), such that their interiors do not intersect (that is, for all $i \neq j$ we have $\Delta_i^\circ \cap \Delta_j^\circ = \emptyset$), and such that their union is equal to Δ . Then $\text{Sub}(\Delta) = \{\Delta_1, \dots, \Delta_k\}$ is called a *Newton subdivision* of Δ .

A *vertex* of the Newton subdivision $\text{Sub}(\Delta)$ is a point in \mathbb{Z}^2 which is a vertex of one of the Δ_i . An *edge* of the Newton subdivision is a line segment in \mathbb{R}^2 which is an edge of some Δ_i .


5.9 Example

The following picture shows two Newton subdivisions, the one on the left is a subdivision of the triangle Δ_4 (see 3.69). The vertices of the Newton subdivisions are drawn in bold. Note that not all integer points in Δ have to be vertices of the Newton subdivision.




5.10 Remark

Not all subdivisions of a Newton polygon are dual to a tropical curve. As an example, consider the following subdivision:



For the edge e in the picture, it is not clear whether we should draw one or two edges dual to it.

As a second example, take the following subdivision:



It is not possible to draw a tropical curve dual to this subdivision. Trying, let us start with the little square in the interior of this subdivision. It would be dual to a crossing of two edges. Next, let us draw the dual of the edge E_3 with direction $(1, -2)$ — the edge e_3 . Then, we draw the dual of E_4 with direction $(1, 1)$ - the edge e_4 . e_4 has to meet e_3 . This is only possible if the edge e_2 is longer than the edge e_1 . But as the same argument works for the other edges adjacent to the dual of the little square, each piece of the edges that cross in the middle has to be longer than the piece counterclockwise next to it. This is of course not possible.

5.11 Definition

A subdivision of a polygon will be called a *regular subdivision* if it is dual to a tropical curve associated to a tropical polynomial as in 2.15.

Equivalently, we can require that the subdivision is induced by a polynomial as its Newton subdivision as in definition 2.13.


Let us determine which Newton subdivisions are dual to unparametrized tropical curves, which occur by our definition as images of simple parametrized tropical curves. Recall that simple parametrized tropical curves are 3-valent, no two vertices are mapped to the same point, no vertex is mapped onto the image of an edge, no edges are mapped to the same line, and through each point pass at most 2 edges. In the dual image, we can therefore only have triangles and parallelograms. (The parallelograms are dual to crossings of two edges.)


5.12 Definition

A regular subdivision $\text{Sub}(\Delta)$ of the Newton polygon Δ is called *simple*, if it contains only triangles and parallelograms.

5.13 Example

The following picture shows a regular simple Newton subdivision. Below, an unparametrized tropical curve dual to it is shown, together with its parametrization.






5.14 Remark

Note that in the picture above the parallelogram is not dual to the image of a vertex of Γ , just to a “crossing” of two edges. If we find a tropical polynomial which defines the same curve as the unparametrized one on the right, then this “crossing” is considered to be a vertex of the tropical curve associated to the tropical polynomial as in chapter 2.

Because our aim here is to describe tropical curves using their dual subdivisions, we will call such a crossing of two edges a *vertex* of the unparametrized tropical curve. (Moreover, all images $h(V)$ of vertices of Γ will be called *vertices* of the unparametrized tropical curve.) Then we also distinguish four different *edges* adjacent to the “crossing-vertex” in the unparametrized tropical curve, even though in the parametrized tropical curve, only two different edges are mapped to this image.



(For all edges of the parametrized tropical curve whose images do not have a common point with another image of an edge, we call the images *edges* of the unparametrized tropical curve.) That is, an unparametrized tropical curve can have more vertices and more edges than its unique simple parametrization. Recall that unparametrized tropical curves are images of simple parametrized tropical curves, therefore this definition of vertices and edges of an unparametrized tropical curve is sufficient.


In this language, theorem 2.15 holds also if we replace “tropical curve associated to a tropical polynomial” by “unparametrized tropical curve” and “the Newton subdivision of F ” by “a Newton subdivision” - each edge of the unparametrized tropical curve is dual to an edge of the Newton subdivision, and each vertex dual to one of the polygons Δ_i of the subdivision. Recall that the integer length of an edge of the Newton subdivision is equal to the weight of the dual edge of the tropical curve.

5.15 Definition

The *degree* of a regular Newton subdivision $\text{Sub}(\Delta)$, $\text{deg}(\text{Sub}(\Delta))$, is defined to be the set of vertices of $\text{Sub}(\Delta)$ which are contained in the boundary of the polygon Δ .

5.16 Example

For the following Newton subdivision, the degree is given by the 10 points $(0, 0)$, $(0, 2)$, $(0, 3)$, $(0, 4)$, $(1, 3)$, $(2, 2)$, $(3, 1)$, $(4, 0)$, $(2, 0)$ and $(1, 0)$. They are marked in bold in the picture.



5.17 Remark

Let us compare this definition with the notions from chapter 4.

There we used the symbol Δ for the degree of a parametrized tropical curve, given by the unordered tuple of directions of the unbounded edges. Each unbounded edge of a simple parametrized tropical curve is of course mapped to an unbounded image in the unparametrized tropical curve. The latter is dual to an edge of the Newton subdivision in the boundary of Δ . The integer length of this edge of the Newton subdivision is equal to the weight of the unbounded edge of the unparametrized tropical curve. If there are no unbounded edges of higher weight, then every integer point on the boundary of Δ is also part of the degree. Then we will by abuse of notation write Δ instead of $\deg(\text{Sub}(\Delta))$. In this case, the number of unbounded edges — denoted $\#\Delta$ in chapter 4 — is equal to the number of integer points in the boundary of the polygon Δ , hence to $\#(\partial\Delta \cap \mathbb{Z}^2)$.

5.18 Remark


Assume we have a tropical curve C given by a tropical polynomial as in definition 2.7. This polynomial defines a Newton polygon, too (see 2.10). This is the same polygon we get when we draw the dual of the unbounded edges of the tropical curve C .

5.19 Example


We want to give an example which illustrates why we choose the triangle Δ_d of 3.69 when we want to count the tropical analogues of degree d curves in \mathbb{P}^2 .

Let C be a curve of degree d in \mathbb{P}_K^2 , where K denotes the completion of the field of Puiseux series as in 2.1. Let C be given by a homogeneous polynomial $f(x, y, z)$ in 3 variables, and assume that the monomials x^d , y^d and z^d appear in the sum f . Restrict C to the open subset $\{z = 1\}$. Then C is given by the polynomial $f(x, y, 1)$. Restrict C furthermore to $\{x \neq 0, y \neq 0\}$, then this restriction can be interpreted as a curve in $(K^*)^2$ and we can consider the image $\overline{\text{Val}(C)}$. By theorem 2.9 the tropical curve $\overline{\text{Val}(C)}$ is equal to the tropical curve associated to the tropical polynomial $\text{trop}(f(x, y, 1))$. Therefore its degree will contain the endpoints $(0, 0)$, $(d, 0)$ and $(0, d)$ of the triangle Δ_d . If we want to count complex projective curves of degree d , we can therefore instead count tropical curves such that the dual Newton polygon is the triangle Δ_d (see remark 4.15).

Let us now come to the genus. To get an idea what the genus of the Newton subdivision should be, consider a loop of an unparametrized tropical curve. It is dual to an interior vertex of the Newton polygon:



However, if the Newton subdivision has a parallelogram next to that interior vertex, we know that this does not come from a 4-valent vertex but from a crossing - hence this is not really a loop then.




These observations justify the following definition:

5.20 Definition

Let Δ be a Newton polygon with a regular simple subdivision $\text{Sub}(\Delta)$. Then the *genus* of this subdivision is defined to be the number of interior vertices of this subdivision minus the number of parallelograms.

5.21 Example

The following pictures show three subdivisions of degree Δ_3 . The one on the left is of genus -1 , because it contains no interior vertex, but a parallelogram. The two on the right are of genus 0 — the one in the middle neither contains interior vertices nor parallelograms. The one on the right has an interior vertex, but also a parallelogram. Below, the dual tropical curves are shown.




Now, we have to define the multiplicity of a Newton subdivision. We will also only define it for simple subdivisions. To get an idea what it should be, recall the definition of

multiplicity from 4.47. We defined first the multiplicity of a 3-valent vertex V . Assume the three flags F_1 , F_2 and F_3 are mapped to V . Then the multiplicity of V is

$$\text{mult}_V = \det(v(F_1), v(F_2)),$$

that is, the area of the parallelogram spanned by the direction vectors of the two flags F_1 and F_2 . (The balancing condition shows that the multiplicity does not depend on the choice of F_1 and F_2 .) Now, dual to the image of the edge $[F_1]$ is a perpendicular edge in the dual Newton subdivision, and the same for $[F_2]$. The area of the parallelogram spanned by the two direction vectors is equal to the area of the parallelogram spanned by the duals of $[F_1]$ and $[F_2]$. But this is the double area of the triangle dual to the vertex V .



Recall that the multiplicity of a 3-valent parametrized tropical curve is defined as the product of the multiplicities of its vertices (which are not adjacent to marked points, but at the moment, we do not consider marked points).

5.22 Definition

Let $\text{Sub}(\Delta)$ be a simple regular subdivision of the Newton polygon Δ . Then the *multiplicity* of this subdivision is defined to be the product of the double areas of all triangles which are contained in the subdivision.

Finally, we have to define irreducibility of a Newton subdivision.

5.23 Remark

Given a Newton subdivision, it does not seem easy to determine whether an unparametrized tropical curve dual to this subdivision is reducible or not (that is, whether it can be parametrized by a disconnected graph Γ or not) without drawing the dual curve. Therefore, we have to use the dual curve to define irreducibility for Newton subdivisions:

5.24 Definition

A regular simple Newton subdivision $\text{Sub}(\Delta)$ is called *irreducible*, if the dual unparametrized tropical curve is irreducible, and *reducible* otherwise.

We now “translated” most of the concepts we introduced for parametrized tropical curves in chapter 4 to regular subdivisions. In the next chapter, we will turn to marked parametrized tropical curves.

5.2. THE DUAL OF A MARKED PARAMETRIZED TROPICAL CURVE

So far we only considered parametrized tropical curves (Γ, h) without marked points, their unparametrized images and dual Newton subdivisions. Now we want to come to marked tropical curves. The basic properties like degree and genus will of course be defined as in the previous section, also in the presence of marked points.

We want to define an analogue of simple parametrized tropical curves also for marked tropical curves. Again, this definition should be made in such a way that the parametrization is unique for a given image $h(\Gamma)$ together with the weights and the images of the marked points on it. We do not want to allow that marked points are mapped to vertices of the unparametrized tropical curve $h(\Gamma)$ — that is, to images of vertices of Γ , respectively to “crossings” of two edges, as they count as vertices of the unparametrized tropical curve, too, by 5.14.

5.25 Definition

A marked parametrized tropical curve $(\Gamma, h, x_1, \dots, x_n)$ is called *simple* if

- Γ is 3-valent,
- h is injective on the set of vertices Γ^0 of Γ ,
- the images $h(V)$ and $h(e)$ of a vertex $V \in \Gamma^0$ and an edge $e \in \Gamma^1 \setminus \{x_1, \dots, x_n\}$ are disjoint, that is $h(V) \cap h(e) = \emptyset$,
- the images $h(e_1)$ and $h(e_2)$ of two edges $e_1, e_2 \in \Gamma^1$, $e_1 \neq e_2$, have at most one point in common, that is $\#\{h(e_1) \cap h(e_2)\} \leq 1$ and
- through each point $p \in \mathbb{R}^2 \setminus \{h(x_1), \dots, h(x_n)\}$ pass at most 2 edges, that is $\#h^{-1}(p) \leq 2$.

5.26 Remark


Note that for a simple marked curve $(\Gamma, h, x_1, \dots, x_n)$ only the marked points x_i can be contracted by h because of the second property of simple curves. Also, no two edges can be mapped to the same line segment, and no vertex to a nonmarked edge. As h is injective on the set of vertices and as there are no 4-valent vertices, marked points cannot be mapped to an image point where three line segments come together. As no vertex can be mapped to a nonmarked edge, marked points cannot be mapped to an intersection point of two edges. Therefore, we can as in 5.2 distinguish the images of the edges and of the vertices. Also, if we think of $h(\Gamma)$ as an unparametrized curve coming from $(\Gamma \setminus \{x_1, \dots, x_n\}, h|_{\Gamma \setminus \{x_1, \dots, x_n\}})$ (where we “straighten the divalent vertices of $\Gamma \setminus \{x_1, \dots, x_n\}$ to get a 3-valent graph again), then the marked points do not lie on the vertices (that is, images of vertices or crossings of two edges) of the unparametrized curve $h(\Gamma)$.

5.27 Definition

Let $(\Gamma, h, x_1, \dots, x_n)$ be a simple marked parametrized curve. Associate the weights $\omega(e)$ for $e \in \Gamma^1 \setminus \{x_1, \dots, x_n\}$ to the images of its edges $h(e)$. Then the image $h(\Gamma)$ together with these weights and together with the images of the marked points $p_i = h(x_i)$ is called a *marked unparametrized tropical curve* $C = (h(\Gamma), \{(h(e), \omega(e))\}, h(x_i))$.

5.28 Example

The following picture shows a marked unparametrized tropical curve.



5.29 Remark

Let C be a marked unparametrized tropical curve.

Recall that C is the image of a simple marked parametrized tropical curve $(\Gamma, h, x_1, \dots, x_n)$.

Now, remove the markings p_i from C , then we get an unparametrized tropical curve C' without markings. By 5.5 we know that it has a unique simple parametrization (Γ_1, h_1) .

As we have seen in remark 5.26, the markings p_i are distinct from the vertices of the (nonmarked) unparametrized tropical curve C' . Therefore we can identify a unique edge e_i of Γ_1 whose image meets p_i . Hence there is only one way to add marked points x_i to the graph Γ_1 in order to get a marked simple parametrization of C .

5.30 Definition

As before we want to specify what we denote by edges and vertices of a marked unparametrized tropical curve. We just remove the markings and consider it as an unparametrized tropical curve - for this, we specified what we denote by edges and vertices in 5.14. More precisely, a *vertex* of the marked unparametrized tropical curve is the image of a vertex of Γ which is not adjacent to a marked point, or a point $h(e_1) \cap h(e_2)$ in the intersection of two images of nonmarked edges $e_1, e_2 \in \Gamma^1 \setminus \{x_1, \dots, x_n\}$. The image of a marked point is *not* called a vertex of the marked unparametrized tropical curve. As before, because we call “crossings” of two edges a vertex, we also distinguish four *edges* adjacent to that vertex. If e_1 and e_2 are two edges of Γ adjacent to a marked point x_i , then the image of e_1, e_2 and x_i is a straight line with the marking $p_i = h(x_i)$ on it. We call this image *one edge* of the marked unparametrized tropical curve, even though it comes from two edges of the graph Γ . The images of all edges whose images do not intersect with the image of another edge and which are not adjacent to marked points are also called *edges* of the marked unparametrized tropical curve.


Now we want to define the dual of a marked unparametrized tropical curve.

5.31 Definition

Let C be a marked unparametrized tropical curve. Forget the markings p_i , then we get an unparametrized curve C' . By remark 5.14 we know that C' is dual to a subdivision $\text{Sub}(\Delta)$. In this subdivision, mark the edges which are dual to edges of C which pass through a marked point p_i .

The set consisting of these *marked edges* is denoted by Ξ .

The picture shows an example. The marked edges Ξ are drawn as thick lines in the dual Newton subdivision.



5.32 Remark

We can consider the Newton subdivision $\text{Sub}(\Delta)$ together with the set Ξ as an analogue of the combinatorial type of a tropical curve.

Given a Newton subdivision, we can choose a marked unparametrized tropical curve C dual to it. (Recall that the weights are equal to the integer lengths of the dual edges.) We can parametrize C by a unique graph (see 5.29). The primitive integral direction of each edge of the graph is prescribed by the Newton subdivision: marked points are mapped with direction 0, all other edges have the directions dual to the Newton subdivision. As we know also the weights of the edges of C , we can determine the directions $v(F)$ for each flag F of Γ . Hence the Newton subdivision together with the marked edges yield a combinatorial type α . The type does not depend on the particular curve C we chose.

Given a (simple) type α , we can take a parametrized tropical curve C in $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$, and the marked unparametrized tropical curve which is defined by it. Then we can draw the dual Newton subdivision and Ξ . The Newton subdivision may however depend on which edges have a point in common in the image.

Now we want to prescribe points $\mathcal{P} = (p_1, \dots, p_n) \in \mathbb{R}^{2n}$ and want to count marked unparametrized tropical curves with markings $\mathcal{P} = (p_1, \dots, p_n)$. As in chapter 4.5, we do not allow all positions for the point configuration \mathcal{P} , but we want them to be in tropical general position (see 4.51). Tropical general position is defined with the aid of the evaluation map: a configuration is called to be in tropical general position, if all preimages in the relevant subset under the evaluation map are 3-valent and of genus g . But when we count the preimages of \mathcal{P} under the evaluation map, we count those curves of a type α where the evaluation map restricted to $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ is not injective with multiplicity 0 due to 4.49. Given a parametrized tropical curve of type α , there is no possibility to check whether the evaluation map restricted to $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ is injective or not at the unparametrized tropical curve $(h(\Gamma), \{(h(e), \omega(e))\}, h(x_i))$. So actually we want to avoid curves of such a type. Also, we want to restrict to simple parametrized tropical curves, because then we can pass to marked unparametrized tropical curves. We therefore make a more restrictive definition of general position for this chapter:

5.33 Definition

A configuration $\mathcal{P} \subset \mathbb{R}^{2n}$ is defined to be in *restricted (tropical) general position* if all relevant preimages under the evaluation map are 3-valent and of genus g , if for all types α that occur in the preimage, the evaluation map restricted to $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ is injective and if all preimages are simple.

Note that the set of points in restricted general position consists still of regions of top dimension: The highest dimension of a stratum $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ that occurs in the relevant subset is $2n$. If the evaluation map is not injective on a stratum, it maps this stratum to a region of at least codimension 1 in \mathbb{R}^{2n} . (See 4.41 and 4.45). If a 3-valent relevant parametrized tropical curve is not simple, then either


- there are two vertices which are mapped to the same point,
- there is a vertex and a nonmarked edge whose images meet in one point,
- there are two edges which are mapped to the same line, or

- there are three edges whose images have a point in common.

Recall that relevant curves do not have contracted bounded edges.

In the first case, the condition that the images of two vertices coincide is a linear condition on the coordinates of $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$.

In the picture, the condition that the images of the two vertices V and V' coincide is a linear condition on the lengths of the four edges e_1 , e_2 , e_3 and e_4 :




Hence, only a lower dimensional subset of $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ fulfills this condition, and this lower dimensional subset is of course also mapped to a region of codimension 1 in \mathbb{R}^{2n} by the linear map ev .

In the second case, it is for example possible that there is a double edge, and a vertex adjacent to one of these edges which is mapped to the same image as the other edge. But as we are in the relevant subset, double edges can only occur if there are 4-valent vertices, so this can actually not happen. Else the condition that the image of a vertex meets the image of a line is also a linear condition on the coordinates of $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ which is only fulfilled by a lower dimensional space. The same is true for the fourth case, too.

In the third case, there are two edges which are mapped to the same line. This is again for example possible if there is a double edge. But as we are in the relevant subset, the curve then has at least one 4-valent vertex.

The condition that two edges which are not linked are mapped to the same line is again a linear condition which is only fulfilled by a lower dimensional space.

For example, in the picture the condition that e_1 and e_2 are mapped to the same line is a condition on the lengths of the edges e_3 and e_4 :



In any case, parametrized tropical curves which are not of genus g , not 3-valent or not simple or where the evaluation map restricted to the corresponding stratum is not injective are mapped to regions of lower dimension in \mathbb{R}^{2n} . Therefore the set of points in restricted general position consists of regions of top dimension separated by these “walls”.

5.34 Lemma

Fix a degree Δ (of parametrized tropical curves, that is, an unordered tuple of directions) and a genus g , and let $n = \#\Delta + g - 1$. Let $\mathcal{P} = (p_1, \dots, p_n)$ be a configuration in restricted general position.

Then the numbers $N_{\text{trop}}(\Delta, g, \mathcal{P})$ of parametrized tropical curves through \mathcal{P} (defined in 4.52) are equal to the numbers of unparametrized tropical curves of genus g and degree Δ with markings p_i counted with multiplicity (as defined in 5.27).

Proof:

Let us first determine the number $N_{\text{trop}}(\Delta, g, \mathcal{P})$. It is equal to the number of preimages of \mathcal{P} under ev in the relevant subset. Recall that $\mathcal{P} = \{p_1, \dots, p_n\}$ is in restricted general position. Hence the set $\text{ev}^{-1}(\mathcal{P})$ is a finite set. This is true because ev is injective on each of the finitely many strata which occur in the preimage. Also, each preimage is simple.

For each preimage, take the corresponding marked unparametrized tropical curve. It is of course of genus g and degree Δ and has the markings p_i . This assignment is a bijection, because by 5.29, there is a unique simple parametrization for a marked unparametrized tropical curve.

Hence we get the same number $N_{\text{trop}}(\Delta, g, \mathcal{P})$ for the unparametrized tropical curves of genus g and degree Δ with markings p_i .

□

5.35 Remark

If Δ is a degree such that the only primitive directions which occur are $(-1, 0)$, $(0, -1)$ and $(1, 1)$, and such that for at least one of these primitive directions all direction vectors have weight 1, then $N_{\text{trop}}(\Delta, g, \mathcal{P})$ is also equal to the number of images $h(\Gamma)$ of parametrized tropical curves of degree Δ and genus g that pass through \mathcal{P} (without the information about the weights as it is needed for unparametrized tropical curves), counted with the multiplicity of the parametrizations.

To see this, take the set of marked unparametrized tropical curves with markings p_i (of which we know by 5.34) that its number — counted with multiplicity — is equal to $N_{\text{trop}}(\Delta, g, \mathcal{P})$. Now forget the information about the weights, then we get a set of images $h(\Gamma)$ of parametrized tropical curves of degree Δ and genus g that pass through \mathcal{P} . We want to show that this assignment is bijective, too. It is injective, because the weights are uniquely determined due to remark 5.7. (This is a consequence of the choice of the degree.)

To see that it is surjective, take the set of images $h(\Gamma)$ of parametrized tropical curves of degree Δ and genus g that pass through \mathcal{P} . Every image of a nonrelevant parametrized curve is counted with multiplicity 0. So the number of these images is equal to the number of images of relevant parametrized tropical curves of degree Δ and genus g that pass through \mathcal{P} . But as \mathcal{P} is in restricted general position, each relevant parametrized tropical curve through \mathcal{P} is simple. Therefore each such image is equal to a marked unparametrized tropical curve of genus g and degree Δ with markings p_i , without the information about the weights.

5.36 Remark

Due to lemma 5.34 we are not going to define the numbers of unparametrized tropical curves through a given configuration \mathcal{P} separately. We know that if the configuration is in restricted general position, this number is just $N_{\text{trop}}(\Delta, g)$ (respectively, $N_{\text{trop}}^{\text{irr}}(\Delta, g)$ if we consider only irreducible curves).

5.3. LATTICE PATHS

The idea how to translate the enumeration of tropical curves through points in general position to Newton subdivisions is to choose in fact a very special (general) position for the points. We will see that the marked edges Ξ dual to a tropical curve through that special point configuration form a lattice path.

5.37 Definition


A path $\gamma : [0, n] \rightarrow \mathbb{R}^2$ is called a *lattice path* if $\gamma|_{[j-1, j]}$, $j = 1, \dots, n$ is an affine-linear map and $\gamma(j) \in \mathbb{Z}^2$ for all $j = 0 \dots, n$.

5.38 Definition

Let λ be a fixed linear map $\lambda : \mathbb{R}^2 \rightarrow \mathbb{R}$ whose kernel has an irrational slope. (For example, $\lambda(x, y) = x - \varepsilon y$, where ε is a small irrational number. This map is in fact the main example we will use in the following.) A lattice path γ is called λ -*increasing* if $\lambda \circ \gamma$ is strictly increasing.

Let p and q be the points in Δ where $\lambda|_{\Delta}$ reaches its minimum (resp. maximum). Then p and q divide the boundary $\partial\Delta$ into two λ -increasing lattice paths $\delta_+ : [0, n_+] \rightarrow \partial\Delta$ (going clockwise around $\partial\Delta$) and $\delta_- : [0, n_-] \rightarrow \partial\Delta$ (going counterclockwise around $\partial\Delta$), where n_{\pm} denotes the number of integer points in the \pm -part of the boundary.

The following picture shows an example for the triangle Δ_3 with vertices $(0, 0)$, $(3, 0)$ and $(0, 3)$ and the map $\lambda(x, y) = x - \varepsilon y$. The image of the path δ_- is drawn as a line, the image of δ_+ as a dotted line.




We will now define the multiplicity of λ -increasing paths as in [23]:

5.39 Definition

Let $\gamma : [0, n] \rightarrow \Delta$ be a λ -increasing path from p to q , that is, $\gamma(0) = p$ and $\gamma(n) = q$. The (*positive and negative*) *multiplicities* $\mu_+(\gamma)$ and $\mu_-(\gamma)$ are defined recursively as follows:

- (1) $\mu_{\pm}(\delta_{\pm}) := 1$.
- (2) If $\gamma \neq \delta_{\pm}$ let $k_{\pm} \in [0, n]$ be the smallest number such that γ makes a left turn (respectively a right turn for μ_-) at $\gamma(k_{\pm})$. (If no such k_{\pm} exists we set $\mu_{\pm}(\gamma) := 0$). Define two other λ -increasing lattice paths γ' and γ'' as follows:

- $\gamma'_\pm : [0, n - 1] \rightarrow \Delta$ is the path that cuts the corner of $\gamma(k_\pm)$, i.e. $\gamma'_\pm(j) := \gamma(j)$ for $j < k_\pm$ and $\gamma'_\pm(j) := \gamma(j + 1)$ for $j \geq k_\pm$.
- $\gamma''_\pm : [0, n] \rightarrow \Delta$ is the path that completes the corner of $\gamma(k_\pm)$ to a parallelogram, i.e. $\gamma''_\pm(j) := \gamma(j)$ for all $j \neq k_\pm$ and $\gamma''_\pm(k_\pm) := \gamma(k_\pm - 1) + \gamma(k_\pm + 1) - \gamma(k_\pm)$:



Let T be the triangle with vertices $\gamma(k_\pm - 1), \gamma(k_\pm), \gamma(k_\pm + 1)$. Then we set

$$\mu_\pm(\gamma) := 2 \cdot \text{Area } T \cdot \mu_\pm(\gamma'_\pm) + \mu_\pm(\gamma''_\pm).$$

As both paths γ'_\pm and γ''_\pm include a smaller area with δ_\pm , we can assume that their multiplicity is known. If γ''_\pm does not map to Δ , $\mu_\pm(\gamma''_\pm)$ is defined to be zero.

Finally, the *multiplicity* $\mu(\gamma)$ is defined to be the product $\mu(\gamma) := \mu_+(\gamma)\mu_-(\gamma)$.


Note that the only end path which does not count zero is the path $\delta_+ : [0, n_+] \rightarrow \Delta$ (respectively δ_-). Paths without a left or right, but not equal to δ_\pm , or “faster” paths $\delta' : [0, n'] \rightarrow \Delta$ such that $\delta_+([0, n_+]) = \delta'([0, n'])$ but $n' < n_+$ have multiplicity zero.

5.40 Remark

Let us interpret the recursion to compute the multiplicity of a path in terms of Newton subdivisions of Δ . Let γ be a λ -increasing path from p to q . Let us first consider the positive multiplicity. The two paths γ and γ' enclose the triangle T . The two paths γ and γ'' enclose a parallelogram. Before we go on to compute the multiplicity of γ' and γ'' recursively, we take two copies of the Newton polygon Δ and mark the triangle in one of them and the parallelogram in the other one. Continuing like this, we get several subdivisions of Δ (above the path γ) in triangles and parallelograms. Of course, we get the same below γ when we perform the recursion for μ_- . Any such subdivision above γ can be combined with any other subdivision below γ . We will call the set of subdivisions which arise like this the *possible Newton subdivisions for γ* . Note that the multiplicity μ of a path γ is nothing else but the number of possible Newton subdivisions for γ counted with their multiplicity as defined in 5.22. Note also that all subdivisions that arise like that contain the edges which are in the image of γ and are simple.

5.41 Example

As an example, the following picture shows a λ -increasing path γ (where $\lambda(x, y) = x - \varepsilon y$ as before) and the two possible Newton subdivisions for γ . Both possible subdivisions are dual to tropical curves of multiplicity 4 (because both have two triangles of size 2). The multiplicity of the path is 8.




5.42 Definition

Let g be an integer and Δ a convex polygon in \mathbb{Z}^2 . Let $e = \#(\partial\Delta \cap \mathbb{Z}^2)$. We denote by $N_{\text{path}}(\Delta, g)$ the number of λ -increasing lattice paths $\gamma : [0, e + g - 1] \rightarrow \Delta$ with $\gamma(0) = p$ and $\gamma(e + g - 1) = q$ counted with their multiplicities as in definition 5.39.

Of course, it is not clear that this definition does not depend on the choice of λ . This will follow from theorem 5.44 which is proved in the next section.

5.43 Example

The following picture shows all λ -increasing paths (where $\lambda(x, y) = x - \varepsilon y$ as before) in Δ_3 with 8 steps. The sum of their multiplicities is $N_{\text{path}}(\Delta_3, 0) = 12$.



5.4. THE CORRESPONDENCE OF TROPICAL CURVES AND LATTICE PATHS

The aim of this section is to prove the following theorem:

5.44 Theorem

For all Δ and g we have $N_{\text{path}}(\Delta, g) = N_{\text{trop}}(\Delta, g)$ (where $N_{\text{path}}(\Delta, g)$ is defined in 5.42 and $N_{\text{trop}}(\Delta, g)$ is defined in 4.73).

See [23] theorem 2.


In fact, we will not only prove that the two numbers coincide. We will choose a certain configuration \mathcal{P}_λ depending on λ , and then show that each possible Newton subdivision for a λ -increasing path (see remark 5.40) is dual to an unparametrized curve passing through \mathcal{P}_λ . So in fact, we give a bijection between the unparametrized tropical curves passing through \mathcal{P}_λ and the set of possible Newton subdivisions for all paths.

A consequence of this proof is that the definition of $N_{\text{path}}(\Delta, g)$ does not depend on the choice of λ : We will show for the point configuration \mathcal{P}_λ depending on λ that the number

of λ -increasing paths is equal to $N_{\text{trop}}(\Delta, g)$. If we choose another map λ' , we also choose a different point configuration $\mathcal{P}_{\lambda'}$. But by theorem 4.53 (respectively, remark 4.73) we know that the number $N_{\text{trop}}(\Delta, g)$ does not depend on the point configuration the curves are required to meet — and therefore we get the same number $N_{\text{path}}(\Delta, g)$ of λ' -increasing paths also for the different choice λ' .

5.45 Remark

We do not state an analogous theorem for the numbers $N_{\text{trop}}^{\text{irr}}(\Delta, g)$. As we have seen in remark 5.23, there is no easy property of a Newton subdivision that decides whether the dual curve is reducible or not. Even more, given a path γ in a Newton polygon Δ , some of the possible Newton subdivisions for it can be dual to reducible curves and others not. Also, given a Newton subdivision which is dual to a reducible curve, there is no better way known to decide which part of the subdivision is dual to which component of the curve than drawing the dual curve.



That is, we could define the numbers $N_{\text{path}}^{\text{irr}}(\Delta, g)$ in the following way: take all λ -increasing paths, determine all possible Newton subdivisions for all paths, draw the dual curves, and count only those which are irreducible. Then this number will obviously coincide with $N_{\text{trop}}^{\text{irr}}(\Delta, g)$. However, this way to determine the number of irreducible possible Newton subdivisions for all λ -increasing paths seems to be a bit complicated. And as we can still not “divide” the Newton subdivisions corresponding to the components of the dual curve, such a definition does not seem very helpful. Indeed, it is possible to reprove the algorithm of Caporaso and Harris (which counts not necessarily irreducible curves) in the dual setting (see chapter 9), but it does not seem so easy to do the same for Kontsevich’s formula (which counts irreducible curves). (Another reason why it seems not easy to reprove Kontsevich’s formula using lattice paths is that we need contracted bounded edges in the tropical proof of Kontsevich’s formula (see chapter 7), which cannot be seen in a dual subdivision.) Therefore we restrict to the numbers $N_{\text{path}}(\Delta, g)$ and do not worry about irreducible curves.

Let us now define the special point configuration we need for the proof of 5.44.

5.46 Definition

Choose a map $\lambda : \mathbb{R}^2 \rightarrow \mathbb{R}$ as in definition 5.38. Choose a line H orthogonal to the kernel of λ and $n = \#(\partial\Delta \cap \mathbb{Z}^2) + g - 1$ points p_1, \dots, p_n on H such that the distance between p_i and p_{i+1} is much bigger than the distance of p_{i-1} and p_i for all i . The choice can in fact be made in such a way that the point configuration is in restricted tropical general

position (see [23], section 8.5). Define the λ -special configuration \mathcal{P}_λ to consist of these n points.


Note that the definition of a tropical curve, of a simple tropical curve and consequently, also the definition of tropical general position is slightly different in [23]. Still, we have that any parametrized tropical curve (as defined here in 4.10) that passes through a configuration of points \mathcal{P} in general position (as defined in [23]) is 3-valent, of genus g and simple, and there is at most one curve of a given type through \mathcal{P} . In fact, what we call “restricted general position” is called “general position” in [23].

5.47 Lemma

Let C be an unparametrized tropical curve through \mathcal{P}_λ as defined in 5.46. Then C intersects the line H (on which the points $\mathcal{P}_\lambda = (p_1, \dots, p_n)$ lie) only in the points p_1, \dots, p_n .

Proof:

Due to the restricted general position of \mathcal{P}_λ , we can conclude that C comes from a unique simple parametrization $(\Gamma, h, x_1, \dots, x_n)$, of type α . Again due to the restricted general position we know that the evaluation map restricted to $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ is injective, therefore by 4.49 the multiplicity of C is nonzero. Hence $(\Gamma, h, x_1, \dots, x_n)$ contains no string. Therefore $\Gamma \setminus \bigcup_i \overline{x_i}$ consists of only rational components with one unbounded end (see 4.50). If C (and hence $h(\Gamma)$) intersects H also in the point p' different from p_1, \dots, p_n , there must be one component K of $\Gamma \setminus \bigcup_i \overline{x_i}$ whose image intersects H in p' .



But then $K \setminus h^{-1}(H)$ consists of two components of which one needs to be compact, as K has only one unbounded end. This is not possible due to the balancing condition. \square

5.48 Lemma

Let C be an unparametrized tropical curve through \mathcal{P}_λ (see 5.46) of genus g and degree Δ . As before, let $\Xi \subset \text{Sub}(\Delta)$ denote the marked edges, which are dual to the edges passing through \mathcal{P}_λ (see definition 5.31).

Then Ξ is the image of a λ -increasing path $\gamma : [0, \#(\partial\Delta \cap \mathbb{Z}^2) + g - 1] \rightarrow \Delta$ from p to q (where p and q are as in definition 5.38 the points of Δ where λ attains its minimum respectively maximum).


Proof:

First note that by lemma 5.47, the set \mathcal{P}_λ coincides with $C \cap H$. That is, we have to see that the edges in $\text{Sub}(\Delta)$ which are dual to edges of C which intersect H form a lattice path.

Consider a vertex V of the Newton subdivision and the edges adjacent to it. Dual to these edges is a chain of edges of C which encloses a convex polyhedron. Any line meets this


chain of edges at most twice. We distinguish several cases depending on the position of V .

- Assume first V is in the interior of Δ .




Then the convex polyhedron is in fact bounded and H meets either none or two of the dual edges. (It cannot meet a vertex, as \mathcal{P}_λ is in restricted tropical general position. Therefore C can be parametrized by a simple marked parametrized tropical curve, and hence none of the points p_i can lie on a vertex of C .) Hence, either none or two marked edges must be adjacent to V .

- Assume next that $V = p$. Recall how p was defined: it is the vertex of Δ where λ attains its minimum. If we draw a line parallel to the kernel of λ through p then this line meets Δ only in p . Even more, the edges adjacent to p in Δ lie on one side of the line parallel to $\ker \lambda$. Assume H intersects two edges of C which are adjacent to p . Change the coordinate system for a moment such that H is of slope 0. Then the slope of one edge must be negative and the slope of the other edge must be positive:



But then the duals of these edges in the Newton polygon Δ would not be on one side of a line parallel to the kernel of λ . So it is not possible that H intersects more than one of the dual edges to the edges adjacent to p .

Assume H intersects none of the dual edges to the interior edges adjacent to p . Then either all those edges lie above H , or below H . Without restriction, assume they lie above H . Then also the unbounded edges dual to the edges in the boundary of Δ adjacent to p lie above H . Also, as both of these edges lie on one side of the line parallel to $\ker \lambda$, one of the dual unbounded edges has to intersect H .



line parallel to $\ker \lambda$

Hence altogether H intersects precisely one of the duals of the edges adjacent to p . So there is one marked edge adjacent to p .

- Assume that $V = q$, then we get the same result as for p : there is one marked edge adjacent to q .
- Assume finally that V is in the boundary of Δ , but neither equal to p nor q .

Then the edges adjacent to V do not lie on one side of a line parallel to $\ker \lambda$ through V . Therefore we can see as before that H intersects two of these edges, if it intersects any at all. So there are either two or no marked edge adjacent to V .

Altogether we have seen that at each vertex V — except p and q — there are either two marked edges adjacent or no marked edge, while at p and q , there is one marked edge adjacent. That is, Ξ is a path from p to q .

At last, we have to see that the path Ξ is λ -increasing. Assume the vertices a_1 , a_2 and a_3 are three consecutive vertices of Ξ , such that the step from a_1 to a_2 is λ -increasing, while the step from a_2 to a_3 is not. But this means, that a_1 and a_3 lie on the same side of a line parallel to $\ker \lambda$ through a_2 . By the above we have seen that then H cannot intersect both dual edges, which contradicts the assumption that the two edges were part of Ξ . \square

Proof of theorem 5.44:

Take the point configuration \mathcal{P}_λ as above in definition 5.46. We know that it is in restricted tropical general position, therefore there are finitely many parametrized tropical curves in $\text{ev}^{-1}(\mathcal{P}_\lambda)$, and each counts with a nonzero multiplicity (see 4.49). Let $(\Gamma, h, x_1, \dots, x_n)$ be one of these curves. Recall that it is simple due to the restricted general position of \mathcal{P}_λ . Then $h(\Gamma)$ is an unparametrized tropical curve of the right genus and degree through \mathcal{P}_λ . If we take the edges of $h(\Gamma)$ that pass through \mathcal{P}_λ and consider their dual edges in the Newton subdivision then these dual edges form a λ -increasing path from p to q by lemma 5.48.

Let γ be a path. We are going to show that there are exactly $\text{mult}(\gamma)$ unparametrized tropical curves (counted with multiplicity) through \mathcal{P}_λ , such that the marked points are dual to $\Xi = \gamma$.

So interpret $\text{Im}(\gamma)$ as a set of marked edges in $\text{Sub}(\Delta)$, and try to draw a dual unparametrized tropical curve. For the edges passing through the points of \mathcal{P}_λ , the direction is prescribed by the path γ and a point through which they should pass is prescribed by \mathcal{P}_λ . Take the “first” marked edge of $\Xi = \text{Im}(\gamma)$ — that is, the one starting at p — and draw a line orthogonal to this marked edge through p_1 . Going on, draw a line dual to the following marked edge of $\Xi = \text{Im}(\gamma)$ through p_2 and so on. We are going to complete these edges to unparametrized tropical curves, counting the possibilities to do this. We will find one unparametrized tropical curve dual to each possible Newton subdivision for γ (see remark 5.40).


The multiplicity μ_+ counts the possibilities to complete the edges dual to γ to a tropical curve (times their multiplicity) in the half-plane above H , whereas μ_- counts below H .

We are going to make this argument precise for μ_+ , for μ_- it is analogous.

Let the first left turn of the path γ be enclosed by the edges E and E' whose duals e and e' pass through p_i and p_{i+1} . The edges through the points p_1, \dots, p_{i-1} do not intersect above H , as this was the first left turn. The edges e and e' will intersect above H , but

below all other possible intersections of dual edges of $\Xi = \text{Im}(\gamma)$. This is true due to the chosen configuration of points: we wanted the distance of p_{j+1} and p_j to be much bigger than the distance of p_j and p_{j-1} . That is, we can draw a parallel line H' to H such that H and H' enclose a strip in which only the intersection point of e and e' lies. Passing from γ to γ' and γ'' corresponds to moving the line H up to H' . The path γ' leaves a triangle T out, and γ'' completes the corner to a parallelogram. These two possibilities are in fact dual to the two possibilities how an unparametrized tropical curve can look like at $e \cap e'$: it can either have a 3-valent vertex — in which case it is dual to the triangle T — or e and e' can just intersect — in which case it is dual to the parallelogram which is enclosed by γ and γ'' . So the change from γ to γ' and γ'' describes the possibilities how a simple unparametrized tropical curve through \mathcal{P}_λ can look like in the strip enclosed by H and H' .

The following picture shows a path γ and the two paths γ' and γ'' , together with the triangle respectively parallelogram which they enclose with γ . Below, the dual curves in the strip enclosed by H and H' are shown.






Recursively, we can see that there is in fact exactly one dual unparametrized tropical curve C through \mathcal{P}_λ to each possible Newton subdivision for γ . C is obviously of degree Δ , as we end up with the two paths δ_+ and δ_- which are dual to the unbounded edges prescribed by Δ in our recursion. (Due to the restricted general position of \mathcal{P}_λ , C is of genus g .)

Hence we constructed a bijection between the set of possible Newton subdivisions for a λ -increasing path and the set of unparametrized tropical curves through \mathcal{P}_λ . The latter number is equal to $N_{\text{trop}}(\Delta, g)$ by lemma 5.34. \square

5.49 Example

The picture shows a λ -increasing path $\gamma : [0, 8] \rightarrow \Delta_3$ (where $\lambda(x, y) = x - \varepsilon y$ as before). Next to it, its two possible Newton subdivisions are shown, and below the dual curves through the point configuration \mathcal{P}_λ . (The distance between the points p_{i+1} and p_i is not

so much bigger than the distance between p_i and p_{i-1} in our picture here, just because this cannot be drawn so easily. The picture should still be sufficient to give an idea on how the special point configuration and the curves through it look like.)



6. THE CORRESPONDENCE OF COMPLEX CURVES AND TROPICAL CURVES

The aim of this chapter is to give a short overview of the proof of Mikhalkin's Correspondence Theorem:

6.1 Theorem (Mikhalkin's Correspondence Theorem)

The numbers $N_{\text{cplx}}^{\text{irr}}(\Delta, g)$ as defined in 3.71 and $N_{\text{trop}}^{\text{irr}}(\Delta, g)$ as defined in 4.52 coincide.

Also, the numbers $N_{\text{cplx}}(\Delta, g)$ from 3.71 and the numbers $N_{\text{trop}}(\Delta, g)$ from 4.73 coincide.

Even more, for a choice \mathcal{P} of points in restricted tropical general position, there exists a configuration $\mathcal{Q} \subset \mathbb{C}^2$ of points in general position (with $\text{Log}(\mathcal{Q}) = \mathcal{P}$) such that for each tropical curve C through \mathcal{P} there are mult C complex curves of genus g and degree Δ through \mathcal{Q} whose amoebas are contained in a neighborhood of C . These curves are distinct for distinct C and irreducible if C is irreducible.

See [23], theorem 1.

In section 3.1 and in section 4.5 we set up enumerative problems both for complex and tropical curves. Theorem 6.1 claims that these two problems coincide — if we want to know one of these numbers, we can determine the other instead. An idea why this should be true can be found in section 2.1: there we chose a different algebraically closed field instead of the complex numbers, the field K which completes the field of Puiseux series. We defined tropical curves as images of such curves over K , and we have seen in theorem 4.27 that the images $h(\Gamma)$ of the parametrized tropical curves which we count in section 4.5 coincide with the closures the images of curves over K under Val (because we only count regular curves). But the numbers $N_{\text{cplx}}^{\text{irr}}(\Delta, g)$ and $N_{\text{cplx}}(\Delta, g)$ do not change if we replace the ground field K by \mathbb{C} .

However, the tropical curves in section 4.5 are counted with a multiplicity (defined in 4.47). We motivated this definition of multiplicity by saying that it should coincide with the number of curves (either of curves in K^2 , or of complex curves) that are mapped to a given tropical curve (either under the valuation map, or by taking logarithm and a limiting process we are about to specify). But the definition we made then for multiplicity was purely combinatorial and it is not obvious that it really coincides with the number of curves (over \mathbb{C} or K). So a main part of the proof of theorem 6.1 is to explain why the multiplicity of a tropical curve coincides with the number of complex curves that map to it.

In section 8.4 we want to generalize Mikhalkin's Correspondence Theorem to tropical curves that satisfy in addition tangency conditions of higher order to a line (a notion which will be defined in section 8.1 and needed for the tropical proof of the algorithm of Caporaso and Harris, given in section 8.2). The knowledge of the idea of the proof of Mikhalkin's Correspondence Theorem is important for our generalization.

6.2 Notation

Let Δ define a toric surface and the class of a curve as in 3.68. Let $g \in \mathbb{Z}$ be given.

During the whole section (x, y) will denote the coordinates of \mathbb{R}^2 , and (z, w) will denote the coordinates of $(\mathbb{C}^*)^2$.

Let \mathcal{Q} be a set of $n = \#(\partial\Delta \cap \mathbb{Z}^2) + g - 1$ points in general position in $(\mathbb{C}^*)^2$ such that $\mathcal{P} = \text{Log}(\mathcal{Q})$ is in restricted tropical general position in \mathbb{R}^2 (defined in 5.33) (where

$$\text{Log} : (\mathbb{C}^*)^2 \rightarrow \mathbb{R}^2 : (z, w) \mapsto (\log |z|, \log |w|)$$

as in chapter 2).

We know that there are finitely many parametrized tropical curves of degree Δ and genus g through \mathcal{P} , and each has a nonzero multiplicity. Call these tropical curves C_1, \dots, C_r . (Note that r need not be equal to $N_{\text{trop}}(\Delta, g)$, and it can differ for different configurations \mathcal{P} .)

Note that by lemma 5.34 it makes no difference whether we consider the curves C_j as parametrized tropical curves or as unparametrized tropical curves.

Mikhalkin's proof is divided into two main parts: In the first part he considers J_t -holomorphic curves, which are in some sense limits of complex curves, and he proves that the amoebas of these limits of complex curves of degree Δ and genus g passing through \mathcal{Q} lie within a small neighborhood of C_j for some j . In the second part he proves that for each C_j there are mult C_j amoebas of J_t -holomorphic curves in a neighborhood of C_j .

Sections 6.2 and 6.3 will each deal with one of these parts of the proof. Section 6.1 starts with some general notions.

6.1. J_t -HOLOMORPHIC CURVES

Recall the notion of an amoeba from the beginning of chapter 2:

6.3 Definition

Let V be a curve in $(\mathbb{C}^*)^2$. Recall that the *amoeba* of V is defined to be the image $\text{Log}(V)$ in \mathbb{R}^2 .

There, we have seen that a tropical curve looks like an amoeba “from far away” — like the limit of an amoeba in some sense. In the following, we will specify this limiting process. We need to shrink the amoeba, until we come from something 2-dimensional to the 1-dimensional “skeleton”. This process of shrinking is done with the map

$$(z, w) \mapsto (\log_t |z|, \log_t |w|)$$

where we let $t \rightarrow \infty$. Note that the image of a given point $(z, w) \in (\mathbb{C}^*)^2$ under this map moves towards the origin if we let $t \rightarrow \infty$. In this sense the map “shrinks” the whole plane in the direction of the origin. Applying this map for a big t to a complex line as in chapter 2 will yield a “thinner” amoeba - something closer to the tropical line. However, we are going to shift this “shrinking” process, we are not going to shrink the real plane, but we are going to shrink the curves in $(\mathbb{C}^*)^2$, before we apply the map Log to get the amoeba in the real plane. This is done by the following map:

6.4 Definition

Let $t > 1$ be a real number. Define

$$H_t : (\mathbb{C}^*)^2 \rightarrow (\mathbb{C}^*)^2 : (z, w) \mapsto \left(|z|^{\frac{1}{\log t}} \frac{z}{|z|}, |w|^{\frac{1}{\log t}} \frac{w}{|w|} \right).$$

A J_t -holomorphic curve V_t is the image $V_t = H_t(V)$ of a holomorphic curve V . The *degree* and *genus* of V_t are defined as the degree and genus of V . A J_t -holomorphic curve V_t is called *irreducible* if its preimage V is irreducible, and *reducible* else.

Note that

$$\begin{aligned} \text{Log} \circ H_t(z, w) &= \text{Log} \left(|z|^{\frac{1}{\log t}} \frac{z}{|z|}, |w|^{\frac{1}{\log t}} \frac{w}{|w|} \right) \\ &= \left(\log |z|^{\frac{1}{\log t}}, \log |w|^{\frac{1}{\log t}} \right) \\ &= \left(\frac{\log |z|}{\log t}, \frac{\log |w|}{\log t} \right) \\ &= \left(\log_t |z|, \log_t |w| \right). \end{aligned}$$

So the map H_t really shifts our shrinking process from above to the curves in $(\mathbb{C}^*)^2$.

What is important about J_t -holomorphic curves is that we can count these instead of “normal” holomorphic curves. This is the result of the following proposition:

6.5 Proposition

For almost all $t > 1$ there are $N_{\text{cplx}}(\Delta, g)$ J_t -holomorphic curves of degree Δ and genus g through \mathcal{Q} .

$N_{\text{cplx}}^{\text{irr}}(\Delta, g)$ of these J_t -holomorphic curves are irreducible.

Proof:

The number of J_t -holomorphic curves through \mathcal{Q} with the required properties is equal to the number of holomorphic curves with the same properties through $H_t^{-1}(\mathcal{Q})$. But $H_t^{-1}(\mathcal{Q})$ is in general position for almost all t .

The statement about the irreducibility follows by definition. □

So to prove theorem 6.1, we need to show that $N_{\text{trop}}(\Delta, g)$ (respectively, $N_{\text{trop}}^{\text{irr}}(\Delta, g)$) is equal to the number of (irreducible) J_t -holomorphic curves of genus g and degree Δ that pass through \mathcal{Q} . This will be shown with the following two lemmata:

6.6 Lemma

Let $\varepsilon > 0$ be given. Then there exists a $T > 1$ such that for all $t > T$ and for all J_t -holomorphic curves V of genus g and degree Δ that pass through \mathcal{Q} , the amoeba $\text{Log}(V)$ is contained in an ε -neighborhood of C_j for some $j = 1, \dots, r$ (where C_j is defined in 6.2).

See lemma 8.3 of [23].

6.7 Lemma

Let $\varepsilon > 0$ be small and $t > 0$ be large. Then the multiplicity $\text{mult } C_j$ as defined in 4.47 is equal to the number of J_t -holomorphic curves V of genus g and degree Δ passing through \mathcal{Q} and such that $\text{Log}(V)$ is contained in the ε -neighborhood of C_j .

Furthermore, if C_j is irreducible any such V is irreducible as well, while if C_j is reducible, any such V is reducible.

See lemma 8.4 of [23]. The following two sections will deal with the proofs of these two lemmata. Using the two statements, we can give a proof of theorem 6.1:

Proof of theorem 6.1:

By proposition 6.5 we know $N_{\text{cplx}}(\Delta, g)$ (respectively, $N_{\text{cplx}}^{\text{irr}}(\Delta, g)$) is equal to the number of J_t -holomorphic curves of genus g and degree Δ that pass through \mathcal{Q} . By lemma 6.6, the amoebas of all J_t -holomorphic curves of genus g and degree Δ that pass through \mathcal{Q} lie each in a small neighborhood of one of the tropical curves C_1, \dots, C_r through \mathcal{P} (see 6.2). For each C_j , there are in fact $\text{mult } C_j$ such J_t -holomorphic curves whose amoeba lie in the neighborhood of C_j by lemma 6.7. As the number $N_{\text{trop}}(\Delta, g)$ is counted with just that multiplicity, the result follows.

The statement about $N_{\text{cplx}}^{\text{irr}}(\Delta, g)$ follows because any irreducible J_t -holomorphic curve projects to an irreducible tropical curve (and the same for a reducible curve) by lemma 6.7. \square

6.2. THE “LIMIT” OF AN AMOEBA — AMOEBAS OF J_t -HOLOMORPHIC CURVES

The aim of this section is to give an idea of the proof of lemma 6.6.

There are two ways of associating a tropical curve to a given J_t -holomorphic curve V — we can associate its tropicalization to it, or the spine of the amoeba $\text{Log}(V)$. Both concepts will be needed to prove that the amoeba of a J_t -holomorphic curve of genus g and degree Δ through \mathcal{Q} lies in a small neighborhood of a C_j . These two tropical curves will be defined in the following:

6.8 Definition

Let V be a (holomorphic) complex curve. It is given by a polynomial $f = \sum_i a_i z^{b_i} w^{c_i}$. Its *tropicalization* $F = \max\{\log|a_i| + b_i \cdot x + c_i \cdot y\}$ defines a tropical curve V^{trop} (see definition 2.7) (we can think of it as an unparametrized tropical curve) that we call the *tropicalization* of V .

Let V_t be a J_t -holomorphic curve. It is then given as the image of a complex curve V under the map H_t of 6.4. We take the tropicalization V^{trop} of V and scale it by the factor $\frac{1}{\log t}$. The scaled tropical curve V_t^{trop} is called the *tropicalization* of V_t .


Note that unlike the situation in section 2.1 (where we dealt with curves over the field K , the completion of the Puiseux series) the tropicalization of V , that is, the tropical curve associated to F , is not equal to the image of V under Log — the latter is a 2-dimensional object, the amoeba.

6.9 Remark

Let V be a J_t -holomorphic curve and $\mathcal{A} = \text{Log}(V)$ its amoeba. By proposition 8.2 of [23], there is a (unparametrized) tropical curve which is contained in the amoeba \mathcal{A} (it can be given by a tropical polynomial whose coefficients depend on the polynomial which defines V). We call this tropical curve the *spine* of the amoeba.

6.10 Example

Here is a picture of an amoeba and its spine:




6.11 Remark


Note that the spine and the tropicalization of a J_t -holomorphic curve can be different. Take for example the complex curve V given by the polynomial

$$\begin{aligned} f(z, w) &= 3z + 3zw + 2z^2 + w^2 + 2w + 1 \\ &= 3z + 3zw + 2z^2 + (w + 1)^2. \end{aligned}$$

Then V has a point of contact order 2 with the line $\{z = 0\}$. Recall the considerations at the beginning of chapter 2: a common point of V and $\{z = 0\}$ leads to a tentacle of the amoeba in direction $x \rightarrow -\infty$. As V has a point of contact order 2, it does not have a second intersection with $\{z = 0\}$. Therefore the amoeba $\text{Log}(V)$ has only one tentacle in this direction. As the spine of $\text{Log}(V)$ is contained in the amoeba, also the spine has only one unbounded edge in direction $(-1, 0)$.



For the tropicalization of V , we can consider the logarithms of the coefficients of the polynomial as “heights” as in 2.13 and project the upper faces of the convex hull of these points to \mathbb{R}^2 to get the dual Newton subdivision of V^{trop} . Doing this, we note that V^{trop} has 2 unbounded edges of direction $(-1, 0)$.



6.12 Notation

Let $(V_k)_{k \in \mathbb{N}}$ be a sequence of curves of genus g and degree Δ passing through \mathcal{Q} such that V_k is a J_{t_k} -holomorphic curve, where $t_k \rightarrow \infty$ for $k \rightarrow \infty$. Let \mathcal{A}_k denote the amoeba $\text{Log}(V_k)$, and S_k the spine of the amoeba $\text{Log}(V_k)$.

Our aim is to see that there is a subsequence of the sequence of the amoebas $(\mathcal{A}_k)_k$ which converges to one of the tropical curves C_j (see 6.2). Both the spines and the tropicalizations of the J_{t_k} -holomorphic curves V_k will be needed to prove this. The following definition is needed to specify the convergence.

6.13 Definition

Let X be a metric space. For two closed subsets $A, B \subset X$ define the *distance of A and B* to be

$$d(A, B) := \max\{ \sup_{a \in A} d(a, B), \sup_{b \in B} d(A, b) \},$$

where $d(a, B)$ denotes the usual distance of a point to a closed set:

$$d(a, B) := \min_{b \in B} \{ d(a, b) \}$$

and $d(a, b)$ denotes the distance.

Let $\{A_k\}$ be a family of closed subsets of X . We say that the family *converges in the Hausdorff metric* to another closed subset A if for every compact set $D \subset X$ there exists a neighborhood U of D such that $\lim_{k \rightarrow \infty} d(A_k \cap U, A \cap U) = 0$.

The following lemma is needed to prove that the amoebas \mathcal{A}_k get “thinner”, the higher k gets.

6.14 Lemma

The amoeba of a J_t -holomorphic curve V is contained in the δ -neighborhood of the tropicalization V^{trop} , where $\delta = \log_t(\#(\Delta \cap \mathbb{Z}^2) - 1)$.

See lemma 8.5 and corollary 8.6 of [23]. Note that V^{trop} is not one of the curves C_j - it does not need to pass through \mathcal{P} , for example.

Idea of the proof:

Assume first V is a holomorphic curve given by the polynomial $f = \sum_i a_i z^{b_i} w^{c_i}$.

Assume a point $(x', y') \in \mathbb{R}^2$ is not contained in the δ -neighborhood of V^{trop} , but in the amoeba $\text{Log}(V)$. Recall that V^{trop} is given as the tropical curve associated to the tropical polynomial $F = \max\{\log |a_i| + b_i \cdot x + c_i \cdot y\}$, where the maximum is taken over all integer points $(b_i, c_i) \in \Delta$. V^{trop} is the set of points where the maximum is attained by at least two terms. If (x', y') is not contained in V^{trop} , the maximum $\max\{\log |a_i| + b_i \cdot x' + c_i \cdot y'\}$ is attained by one single term. Assume it is attained by the term $\log |a_{i'}| + b_{i'} x' + c_{i'} y'$. Because we assume that (x', y') is not even contained in the δ -neighborhood of V^{trop} , we can in fact conclude that

$$\log |a_{i'}| + b_{i'} x' + c_{i'} y' > \log |a_i| + b_i \cdot x' + c_i \cdot y' + \delta$$

for all $i \neq i'$.

But (x', y') is contained in $\text{Log}(V)$, therefore there is a point $(z', w') \in V$ such that $\text{Log}(z', w') = (x', y')$. As $(z', w') \in V$, we have $f(z', w') = 0$. The triangle inequality implies

$$|a_{i'} z'^{b_{i'}} w'^{c_{i'}}| \leq \left| \sum_{i \neq i'} a_i z'^{b_i} w'^{c_i} \right|.$$

We apply \log to both sides of this inequality and get

$$\begin{aligned} & \log |a_{i'}| + b_{i'} x' + c_{i'} y' \\ &= \log |a_{i'} z'^{b_{i'}} w'^{c_{i'}}| \\ &\leq \log \left| \sum_{i \neq i'} a_i z'^{b_i} w'^{c_i} \right| \\ &\leq \log \left((\#(\Delta \cap \mathbb{Z}^2) - 1) \cdot \max_{i \neq i'} \{|a_i z'^{b_i} w'^{c_i}|\} \right) \\ &= \delta + \max_{i \neq i'} \{\log |a_i| + b_i x' + c_i y'\}. \end{aligned}$$

Hence we get a contradiction to the above.

The case of a J_t -holomorphic curve follows from the result for the holomorphic preimage under H_t . \square

6.15 Proposition

Take the sequence $(V_k)_k$ of J_{t_k} -holomorphic curves from 6.12. Let $(S_k)_k$ denote the corresponding sequence of spines. Then there is a subsequence of $(S_k)_k$ which converges in the Hausdorff metric to one of the tropical curves C_j (see 6.2).

See proposition 8.7 of [23].

Idea of the proof:

Recall that the spine of an amoeba can be given by a tropical polynomial which depends on the amoeba. Hence for the sequence $(S_k)_k$ of spines, we also get a sequence of tropical polynomials describing the spines. Then there is a subsequence which converges to a tropical curve which is given by the “limit” of these tropical polynomials — for a detailed description of this limit (there is some ambiguity in the coefficients of the polynomials) and for a proof of this statement, see proposition 3.9 of [23].

It remains to show that the limit of the converging subsequence is a tropical curve of genus g and degree Δ that passes through \mathcal{P} — hence one of the curves C_j . This is done in proposition 8.7 of [23]. \square

Idea of the proof of 6.6:

Take again the sequence $(V_k)_k$ of J_{t_k} -holomorphic curves from 6.12. Corollary 8.8 of [23] shows that the corresponding sequence of spines $(S_k)_k$ is a union of subsequences, each of which consists either of finitely many terms, or converges to one of the curves C_j from 6.2.

Lemma 6.14 is needed to prove that not only the spines S_k , but also the amoebas \mathcal{A}_k converge. Of course, the amoeba surrounds the spine, and if the spines converge to C_j , then the amoeba has to be close to C_j , too. However, the amoeba is a 2-dimensional object which has a “thickness”, and just that S_k is close to C_j does not imply that \mathcal{A}_k

is indeed contained in a small neighborhood. But in 6.14 we have seen that it is in the δ -neighborhood of another tropical curve, of the tropicalization of V_k . In particular this means that its thickness cannot be bigger than 2δ . The number δ depends on t_k , for the J_{t_k} -holomorphic curve V_k it is $\delta = \log_{t_k}(\#(\Delta \cap \mathbb{Z}^2) - 1)$. As $t_k \rightarrow \infty$ when $k \rightarrow \infty$, the thickness of the amoebas \mathcal{A}_k gets smaller, the higher k gets. In particular, \mathcal{A}_k is contained in a small neighborhood of C_j .

Note that a spine S_k which is contained in the sequence that converges to C_j can have more edges than C_j , but these edges vanish in the limit. All other edges of S_k tend to a parallel edge of C_j . More precisely, proposition 8.9 of [23] yields a small value $\tilde{\delta}(t_k)$ depending on t_k such that all edges of the spine S_k are in a $\tilde{\delta}(t_k)$ -neighborhood of the corresponding edge of C_j . (The proof of 8.9 of [23] uses — among others — lemma 6.14 again: the spine S_k is contained in \mathcal{A}_k . We know $\mathcal{A}_k = \text{Log}(V_k)$ contains the points \mathcal{P} , because V_k contains the points \mathcal{Q} . An edge of S_k which corresponds to an edge e of C that passes through some p_i cannot have bigger distance from e than δ , because both p_i and S_k are contained in \mathcal{A}_k .) As furthermore by 6.14, the thickness of the amoeba \mathcal{A}_k is smaller than $\delta = \log_{t_k}(\#(\Delta \cap \mathbb{Z}^2) - 1)$, we can conclude that \mathcal{A}_k is contained in a $\tilde{\delta}(t_k) + 2\delta$ -neighborhood of C_j . \square

6.3. THE NUMBER OF COMPLEX CURVES WHOSE LIMIT IS A GIVEN TROPICAL CURVE

The aim of this section is to give an overview of the proof of the lemma 6.7.

For the whole section, let $C = C_j$ be one of the tropical curves through \mathcal{P} as in 6.2.

The proof works in two steps — we are going to define complex tropical curves (they can be thought of as J_∞ -holomorphic curves) and show how many complex tropical curves project to C under Log in 6.22. Then we are going to show how many J_t -holomorphic curves are contained in a neighborhood of each such complex tropical curve that projects to C in 6.24 and 6.25.

6.16 Definition

Let $(V_k)_k$ be a sequence of J_{t_k} -holomorphic curves, where $t_k \rightarrow \infty$ for $k \rightarrow \infty$. Assume the sequence converges in the Hausdorff metric to V_∞ . Then V_∞ is called a *complex tropical curve*.

V_∞ is defined to be of *genus* g , if it is the limit of a sequence of J_{t_k} -holomorphic curves of genus g and cannot be presented as the limit of a sequence of J_{t_k} -holomorphic curves of smaller genus. The *degree* of V_∞ is defined to be the degree of the tropical curve $\text{Log}(V_\infty)$.

There are other ways to characterize a complex tropical curve, see proposition 6.1 of [23].

The definition includes the claim that $\text{Log}(V_\infty)$ is in fact a tropical curve. This follows from one of the other characterizations of complex tropical curves (see definition 6.17 below) and Kapranov's theorem (see proposition 6.3 of [23] and 2.9). We are going to present this different characterization of a complex tropical curve, and leave the proof that it coincides with definition 6.16 to proposition 6.1 of [23].

Recall the completion K of the field of Puiseux series defined in section 2.1 with the valuation $\text{val} : K \rightarrow \mathbb{R} \cup \{-\infty\}$. To a Puiseux series

$$p(t) = a_1 t^{q_1} + a_2 t^{q_2} + a_3 t^{q_3} + \dots$$

(where $a_i \in \mathbb{C}$ and $q_1 < q_2 \dots$) the valuation associates the rational number $-q_1$.

This valuation can be “complexified” to give a map

$$w : p(t) \mapsto e^{-q_1 + i \cdot \arg(a_1)}.$$

In fact, the valuation can be complexified to w on the whole completion K , not only for the Puiseux series.

Applying this map componentwise, we get a map $W : (K^*)^2 \rightarrow (\mathbb{C}^*)^2$ such that $\text{Log} \circ W = \text{Val}$.

6.17 Definition

Let $V \subset (K^*)^2$ be a curve. Then the image $W(V) \subset (\mathbb{C}^*)^2$ is defined to be a *complex tropical curve*.

6.18 Definition

Let $C = C_j$ be one of the tropical curves through \mathcal{P} from 6.2. Think of $C = (\Gamma, h, x_1, \dots, x_n)$ as a parametrized tropical curve. Each edge of Γ has a certain weight which is equal to the factor with which we need to multiply the primitive integer vector u to get the direction v (see remark 4.13). We define the *edge multiplicity* of C and \mathcal{P} $\mu_{\text{edge}}(C, \mathcal{P})$ to be the product of all those weights.


6.19 Remark

Note that it is important for definition 6.18 that we think of C as a parametrized tropical curve. Each marked point is adjacent to two other edges of Γ which are mapped to the same edge of the unparametrized tropical curve as in 5.30. Therefore, when computing the edge multiplicity of C , we multiply twice with the weight of the direction of this line.

If we think of C as an unparametrized tropical curve, we have to multiply with the square of the weights of those edges which pass through a point of \mathcal{P} instead. This justifies the notation $\mu_{\text{edge}}(C, \mathcal{P})$ — considered as an unparametrized tropical curve, the edge multiplicity depends not only on C but also on the position of the points \mathcal{P} on C .

6.20 Example

The edge multiplicity of the following tropical curve through the points p_1, p_2 and p_3 is $2 \cdot 2 = 4$:



6.21 Proposition

Let Δ' be a lattice triangle. Let q_1 and q_2 in $(\mathbb{C}^*)^2$ be two points in general position,

such that $p_1 = \text{Log}(q_1)$ and $p_2 = \text{Log}(q_2)$ are in restricted tropical general position (with respect to degree Δ' and genus 0). There is exactly one rational tropical curve C (with only one vertex) dual to Δ' through p_1 and p_2 . Assume the two edges that pass through p_1 respectively p_2 are of weight ω_1 respectively ω_2 .


Then there are

$$\frac{2 \text{Area}(\Delta')}{\omega_1 \omega_2}$$

rational complex tropical curves that pass through q_1 and q_2 and project to C under Log .

For a proof, see proposition 6.17 of [23].

The following picture shows an example. The area of Δ' is 3. The weights of the two edges that pass through the two points are 2 and 1. So the proposition claims that there are 3 complex tropical curves through q_1 and q_2 that project to this tropical curve under Log .



6.22 Proposition

Let $C = C_j$ be one of the tropical curves through \mathcal{P} as in 6.2.

Then there are


$$\frac{\text{mult } C}{\mu_{\text{edge}}(C, \mathcal{P})}$$

complex tropical curves in $(\mathbb{C}^*)^2$ of genus g and degree Δ that pass through \mathcal{Q} and project to C under Log .

See proposition 6.18 of [23].

Proof:

As \mathcal{P} is in restricted tropical general position, we know that C contains no string (4.49). Think of $C = (\Gamma, h, x_1, \dots, x_n)$ again as a parametrized tropical curve. Then by 4.50 $\Gamma \setminus \bigcup_i \overline{x_i}$ consists of only rational components each of which contains precisely one unbounded edge.



Let K be such a component. Let x_i and x_j be two marked points which are adjacent to two edges of K that are adjacent to the same 3-valent vertex V . Let p_i and p_j in \mathcal{P} be the

two points to which the two marked points x_i and x_j are mapped. Let Δ' be the triangle dual to the vertex V . Then there are $2 \cdot \text{Area}(\Delta')/(\omega_1\omega_2)$ complex tropical curves that pass through q_i and q_j and project to the dual of Δ' . The result follows by induction. \square

To prove lemma 6.7, it remains to show that there are $\mu_{\text{edge}}(C, \mathcal{P})$ J_t -holomorphic curves of genus g and degree Δ passing through \mathcal{Q} in the neighborhood of each complex tropical curve which maps to C under Log . This is the hardest part of the proof of theorem 6.1. We will only give a short overview of the proof. The main idea is to prove the statement separately for each polygon in the Newton subdivision of Δ dual to C — for each edge, triangle and parallelogram — and to “glue” the J_t -holomorphic curves corresponding to each such polygon using Viro’s Patchworking method (see proposition 8.12 of [23]).

C can be considered as an unparametrized tropical curve. Also, we can find a tropical polynomial which defines C . (That this is indeed possible is a consequence of theorem 4.27: C comes from a curve over the completion K of the field of Puiseux series, which is given by a polynomial f . Then $\text{trop}(f)$ is a suitable tropical polynomial which defines C (see 2.9).)

So we can think of C as given by a tropical polynomial $F = \max\{a_i + b_i x + c_i y\}$, where the sum is taken over all i such that (b_i, c_i) is an integer point in the polygon Δ . There is in fact some ambiguity in the choice of the coefficients a_i — for more details, we refer to [23].

Now let V_∞ be one of the complex tropical curves which maps to C under Log . By the second characterization of complex tropical curves (see definition 6.17) we know that $V_\infty = W(V)$, where V is a curve over the completion of the field of Puiseux series K . V is given by a polynomial over K . The coefficients a'_i of this polynomial can be mapped to $(\mathbb{C}^*)^2$ by the map w . These images $w(a'_i)$ are then in fact determined (up to some ambiguity) by the complex tropical curve V_∞ to which V is mapped via W .

That is, given C and given a complex tropical curve V_∞ that maps to C , we get two sets of coefficients: a set of real coefficients a_i (for the tropical polynomial defining C) and a set of complex coefficients $w(a'_i)$ (for the images under w of the coefficients a'_i of the polynomial which defines the curve V over K that is mapped to V_∞ under W).

Let $f_t^\zeta = \sum \arg(\zeta_i) t^{\log \zeta_i} z^{b_i} w^{c_i}$ (where the sum is again taken over all i such that (b_i, c_i) is an integer point in the polygon Δ) be a polynomial depending on t and a set of complex coefficients ζ .

Now assume V_t is a J_t -holomorphic curve in the neighborhood of V_∞ .

Then $V_t = V_t^\zeta$ can be presented as $V_t^\zeta = H_t(\{f_t^\zeta = 0\})$, and the set of coefficients ζ has to satisfy a condition given by the two sets of coefficients a_i and $w(a'_i)$ we just described. (For a proof of this statement and for more detailed descriptions on the conditions and the two sets of coefficients, see proposition 8.11 of [23].)

The aim is now to count those J_t -holomorphic curves V_t^ζ in the neighborhood of V_∞ which are suitable — that is, which are of the right genus and degree, and pass through \mathcal{Q} . We call a set ζ of complex coefficients which define such a suitable J_t -holomorphic curve a *suitable set* ζ .

To count the suitable sets ζ , we work with a version of Viro's patchworking method.

We cover the tropical curve C by small neighborhoods around the vertices, edges and crossings of two edges. Each vertex, edge or crossing of two edges is dual to a polygon Δ' — edge, triangle or parallelogram — of the Newton subdivision $\text{Sub}(\Delta)$ dual to C .

The patchworking principle roughly states now that if we change the coefficients ζ_i corresponding to points (b_i, c_i) which *do not* belong to the polygon Δ' of $\text{Sub}(\Delta)$, then this change has little effect on the curve V_t^ζ restricted to the open subset around the edge, vertex or crossing dual to Δ' (see proposition 8.12 of [23]).

Proposition 8.14 of [23] describes how the curves V_t^ζ restricted to the open subset around the dual of each polygon Δ' have to look like topologically in order to have a curve of genus g in total.

6.23 Remark

A corollary of this proposition (corollary 8.15 of [23]) is then that each such curve V_t^ζ of genus g which is irreducible maps to an irreducible tropical curve C , and each reducible V_t^ζ maps to a reducible C .

The lemmata 8.16, 8.17 and 8.21 of [23] count for each polygon Δ' of $\text{Sub}(\Delta)$ the numbers of suitable coefficients ζ (which satisfy that the curve V_t^ζ — restricted to the open subset around the dual of Δ' — fulfills the conditions of 8.14, and passes through \mathcal{Q}). That is, these lemmata count the numbers of suitable coefficients ζ for each polygon separately, for which the “glued” curve, V_t^ζ , is indeed a suitable curve.

However, integer points of Δ may belong to several polygons — edges, triangles or parallelograms — of $\text{Sub}(\Delta)$. So it is not clear if we can choose the suitable ζ for each Δ' separately. Therefore, we have to choose an order on the polygons Δ' . Then we choose the coefficients ζ for each Δ' one after the other following this order.

It is still not easy to show how many possibilities there are to choose the new coefficients in each step such that they are compatible to the old coefficients. In fact, the following proposition (see 8.23 of [23]) shows this only for the case that $\text{Sub}(\Delta)$ contains no edge of a higher integer length. That is, C contains no edge of a higher weight, and the edge multiplicity of C and \mathcal{P} is 1.

6.24 Proposition

Let C be a tropical curve through \mathcal{P} such that C contains no edge of weight greater 1. Let V_∞ be a complex tropical curve that maps to C . Then there is one suitable choice for ζ such that V_t^ζ is a J_t -holomorphic curve of degree Δ and genus g passing through \mathcal{Q} .

In particular, there is $1 = \mu_{\text{edge}}(C, \mathcal{P})$ J_t -holomorphic curve of degree Δ and genus g passing through \mathcal{Q} in the neighborhood of V_∞ .

For a proof, see proposition 8.23 of [23]. It uses the three lemmata 8.16, 8.17 and 8.21 which count the suitable coefficients separately for each Δ' , and the chosen order on the Δ' that guarantees that we make compatible choices.

It remains to prove the analogous statement for the case that C contains edges with higher weights. As we chose $n = \#(\partial\Delta \cap \mathbb{Z}^2) + g - 1$ points \mathcal{P} , the unbounded edges of C all need

to have weight 1. So only an interior edge can have a higher weight. Assume e is an edge (of the unparametrized tropical curve C) of weight ω . Assume e is dual to the edge Δ' of the dual Newton subdivision of C which then has integer length ω . If e is disjoint from the points \mathcal{P} , it contributes ω to the edge multiplicity, else, it contributes ω^2 . So in order to see that there are $\mu_{\text{edge}}(C, \mathcal{P})$ J_t -holomorphic curves contained in the neighborhood of each complex tropical curve which projects to C , we have to show that there are ω choices for suitable coefficients for Δ' if e is disjoint from \mathcal{P} , and ω^2 else.

6.25 Lemma

Let C be a tropical curve through \mathcal{P} with a bounded edge e of weight $\omega > 1$. Let e be dual to the edge Δ' of the dual Newton subdivision of C .

Then we have

- ω suitable choices for the coefficients ζ of Δ' if e does not pass through one of the points of \mathcal{P} (that is, if $\Delta' \not\subset \Xi$), and
- ω^2 suitable choices if e passes through one of the points of \mathcal{P} (that is, if $\Delta' \subset \Xi$).


Recall that suitable means in this context: compatible with the choices we already made for polygons which occurred earlier in our order, passing through \mathcal{Q} , and such that V_t^ζ restricted to the open subset around e is of the right form to guarantee that the complete curve V_t^ζ is of genus g (where the “right form” is specified in proposition 8.14 of [23]). For more details see lemma 8.24 of [23].

Idea of the proof:

By the patchworking principle we can conclude that it is not important which Newton polygon Δ surrounds the edge Δ' of integer length ω . Therefore we can assume that Δ' lies in a special polygon.

We are also going to assume that $\Delta' \not\subset \Xi$ and that ω is odd. The other cases work analogously.

So we can assume Δ' lies in the parallelogram with the vertices $(0, 0)$, $(0, -1)$, $(\omega, 0)$ and $(\omega, 1)$.



Δ


The tropical curve C dual to this Newton subdivision has 4 unbounded edges and is rational — therefore it is fixed by 3 points p_1 , p_2 and p_3 . We can choose p_1 , p_2 and p_3 so that there is only one rational tropical curve of degree Δ passing through the three points — the curve C — and such that Ξ is equal to three edges of the boundary:



Δ

Also, choose preimages q_1, q_2 and q_3 under Log in general position.

We can furthermore choose a different set of points p'_1, p'_2 and p'_3 (and preimages q'_1, q'_2 and q'_3) such that the only tropical curve C' of degree Δ and genus 0 that passes through p'_1, p'_2 and p'_3 has the following Newton subdivision (and the same marked edges Ξ):



Note that C' has no edge of a higher weight. We know $\text{mult } C' = \omega^2$ and $\mu_{\text{edge}}(C', \{p'_1, p'_2, p'_3\}) = 1$. Therefore by 6.22, there are ω^2 complex tropical curves that project to C' .

As C' has no edge of a higher weight we can apply proposition 6.24 and see that there is one suitable J_t -holomorphic curve in the neighborhood of each of the ω^2 complex tropical curves.

We have chosen the point configuration such that C' is the only rational tropical curve of degree Δ that passes through it. Therefore, the amoebas of all rational J_t -holomorphic curves of degree Δ passing through q'_1, q'_2 and q'_3 have to lie in the neighborhood of C' by 6.6. That is, $N_{\text{cplx}}(\Delta, 0)$ (respectively $N_{\text{cplx}}^{\text{irr}}(\Delta, 0)$) is equal to ω^2 .

For C , we have $\text{mult } C = \omega^2$ and $\mu_{\text{edge}}(C, \{p_1, p_2, p_3\}) = \omega$, therefore by 6.22, there are ω complex tropical curves that project to C .

Again, because C is the only rational tropical curve of degree Δ that passes through p_1, p_2 and p_3 , every rational J_t -holomorphic curve of degree Δ passing through q_1, q_2 and q_3 has to map to a neighborhood of C under Log .

Therefore, there have to be altogether $N_{\text{cplx}}(\Delta, 0)$ J_t -holomorphic curves in the neighborhoods of these ω complex tropical curves. We computed this number with the help of C' , it is ω^2 . By symmetry, there have to be ω J_t -holomorphic curves in the neighborhood of each complex tropical curve that projects to C . \square

6.26 Remark

Note that for this proof it is indeed necessary that e is a bounded edge of C . The proof cannot be generalized to the case of a tropical curve which has unbounded edges of a higher weight.

6.27 Remark

Another way of explaining the result of 6.25 is the following:

Let e again be an edge of C of weight ω , and Δ' the dual edge in $\text{Sub}(\Delta)$. Let A and B be the two vertices adjacent to e , and Δ_A and Δ_B be the two triangles in $\text{Sub}(\Delta)$ dual to A and B .

As in remark 3.68, the triangles Δ_A and Δ_B describe a toric surface, and the edge Δ' which belongs to both triangles (respectively, another edge of integer length ω if Δ' belongs to a parallelogram) defines a divisor $D_{\Delta'}$ on each toric surface.

The suitable curves V_t^ζ restricted to the neighborhood around A (respectively B) approximate curves which have tangency order ω with the divisor $D_{\Delta'}$ in the respective toric surface (see [23], remark 8.25). The curves cannot have more intersection points with this divisor, as otherwise the genus of the whole curve V_t^ζ (after gluing each little piece) would be too big.

We can think of the two curves on the toric surfaces for Δ_A and Δ_B as two annuli which are circled almost around themselves ω times. There are ω ways to identify them. If there is in addition a point q_i , there are also ω choices for the sheet on which q_i lies.

Let us sum up the contents of this section to prove lemma 6.7:

Idea of the proof of lemma 6.7:

Let C be one of the tropical curves through \mathcal{P} (as in 6.2). By proposition 6.22 we know that there are $\text{mult } C / \mu_{\text{edge}}(C, \mathcal{P})$ complex tropical curves which project to C under Log . By 6.24 and 6.25 we know that there are $\mu_{\text{edge}}(C, \mathcal{P})$ J_t -holomorphic curves of genus g and degree Δ through \mathcal{Q} in a neighborhood of each such complex tropical curve. Altogether we have $\text{mult } C$ J_t -holomorphic curves of genus g and degree Δ through \mathcal{Q} whose amoebas lie in a neighborhood of C .

The statement about the irreducibility follows due to remark 6.23. □

7. A TROPICAL PROOF OF KONTSEVICH'S FORMULA

As we have seen in chapter 6, the numbers $N_{\text{cplx}}^{\text{irr}}(d, g)$ and $N_{\text{trop}}^{\text{irr}}(d, g)$ coincide. For rational curves, we know an algorithm — Kontsevich's formula (see section 3.2) — which allows to compute the numbers $N_{\text{cplx}}^{\text{irr}}(d, 0)$ recursively. Hence we can conclude that the tropical numbers $N_{\text{trop}}^{\text{irr}}(d, 0)$ must fulfill the same recursion formula. The aim of this section is to find a reason in the tropical world why the numbers $N_{\text{trop}}^{\text{irr}}(d, 0)$ satisfy Kontsevich's formula. More precisely, the aim of this chapter is to prove the following theorem tropically:

7.1 Theorem

The numbers $N_{\text{trop}}^{\text{irr}}(d, 0)$ (defined in 4.52) satisfy Kontsevich's formula (as defined in 3.30).

Our tropical proof of Kontsevich's formula contributes to the “translation” of complex geometry to tropical geometry. Furthermore, it shows that the two fields are very much related: we will see that the tropical proof follows essentially the same ideas as the classical one (see section 3.2).

7.2 Remark

Using theorem 7.1 and theorem 3.31 we can in fact give an alternative proof of Mikhalkin's Correspondence Theorem 6.1 for rational curves: knowing that both the numbers $N_{\text{cplx}}^{\text{irr}}(d, 0)$ and $N_{\text{trop}}^{\text{irr}}(d, 0)$ satisfy the same recursion formula, and knowing that there is one tropical as well as one complex line through two points, we can conclude recursively that $N_{\text{cplx}}^{\text{irr}}(d, 0) = N_{\text{trop}}^{\text{irr}}(d, 0)$.

In section 7.1, we reconsider the moduli space of tropical curves. In the case of rational curves, we can in fact choose an easier definition of the moduli space than the one we used in chapter 4: we do not need to pass to the relevant subset. In section 7.2, we define tropical analogues of an important tool used in the classical proof: forgetful maps. The most important forgetful map forgets the map h and all marked points but the first four. (If we do not specify in the following which forgetful map we are talking about, this is the one.) We cannot work with divisors and intersection theory as these concepts are not yet fully developed in the tropical world. Therefore, we have to replace the statement that the two boundary divisors $D_{1,2/3,4}$ and $D_{1,3/2,4}$ in the moduli space of stable maps are linearly equivalent somehow in the tropical world. To do this, we combine the forgetful map with the evaluation map: we impose the incidence conditions we need for Kontsevich's formula (that is, we will require the tropical curves to meet two lines and $3d - 2$ points just as in the classical proof) and we require that the curves map to a given point in $\mathcal{M}_{\text{trop}, 0,4}$ under the forgetful map. We show in proposition 7.14 that the number of curves satisfying these conditions does not depend on the special choice of the point in $\mathcal{M}_{\text{trop}, 0,4}$. The idea to prove this is the same as for theorem 4.53 which shows that the number of tropical curves passing through a given set of points, $N_{\text{trop}}^{\text{irr}}(\Delta, g)$, does not depend on the special choice of points (although the multiplicity with which the curves have to be counted is not the same).

In section 7.3 we will choose two special points in $\mathcal{M}_{\text{trop}, 0,4}$ — namely two points where the length of the bounded edge which links the four markings x_1, \dots, x_4 is very large.

Using proposition 7.14 we know that the numbers of tropical curves which satisfy the incidence conditions and map to the two special points in $\mathcal{M}_{\text{trop}, 0, 4}$ are equal for both points. Finally, we have to interpret these two numbers in terms of reducible curves consisting of two components of lower degree. We will see that the two numbers can be interpreted as sums just analogously to the classical case.

The tropical proof of Kontsevich's formula shows that it is possible to carry many concepts from classical complex geometry over to the tropical world: moduli spaces, morphisms, divisors and divisor classes, intersection multiplicities, and so on. Even if we only make these constructions in the specific cases we need, we hope that our work will be useful to find the correct definitions of these concepts in the general tropical setting.

Note that it is not straightforward to generalize the methods used in this chapter to curves of arbitrary degree: the special degree is needed in proposition 7.15.

The result we describe in this chapter was achieved in joint work with Andreas Gathmann and published as preprint in [13].

7.1. THE ENUMERATIVE PROBLEM FOR RATIONAL PARAMETRIZED TROPICAL CURVES

In chapter 4 we decided to work with the subset of relevant tropical curves (see definition 4.36). The reason is that restricting to relevant curves we were able to give a good bound on the dimensions of the strata $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ (see proposition 4.41). We needed the knowledge about these dimensions in order to prove the main theorem of chapter 4, theorem 4.53. We will see that for rational curves, this reason is absent.

7.3 Lemma

A rational parametrized tropical curve C is regular, that is, the dimension of the stratum $\mathcal{M}_{\text{trop}, 0, n}^\alpha(\Delta)$ (where α denotes the combinatorial type of C) is equal to the expected dimension (defined in 4.24).

Proof:

Recall that by lemma 4.21 the strata $\mathcal{M}_{\text{trop}, g, n}^\alpha(\Delta)$ are given as the subset of $\mathbb{R}^{2+\#\Gamma_0^1}$ where the coordinates given by the lengths are positive and the conditions that the loops close up are satisfied. For rational curves, there are of course no loops and hence we have

$$\dim \mathcal{M}_{\text{trop}, 0, n}^\alpha(\Delta) = 2 + \#\Gamma_0^1 = \text{edim} \mathcal{M}_{\text{trop}, 0, n}^\alpha(\Delta).$$

□

In particular, we can determine the dimensions of all strata $\mathcal{M}_{\text{trop}, 0, n}^\alpha(\Delta)$.

So the reason for which we passed to the relevant subset in chapter 4 has become absent when we work with rational curves. Therefore, we want to work with the space $\mathcal{M}_{\text{trop}, 0, n}(\Delta)$ here instead of $\widetilde{\mathcal{M}}_{\text{trop}, 0, n}(\Delta)$. This is important, because we need to consider curves with contracted bounded edges later on.

7.4 Notation

For the whole chapter, we work with the space $\mathcal{M}_{\text{trop}, 0, n}(\Delta)$ as moduli space of tropical curves. We denote by

$$\text{ev} : \mathcal{M}_{\text{trop}, 0, n}(\Delta) \rightarrow \mathbb{R}^{2n} : (\Gamma, h, x_1, \dots, x_n) \mapsto (h(x_1), \dots, h(x_n))$$

the evaluation map now starting from this space.

Analogously to 4.45, the evaluation map is a linear map restricted to each stratum $\mathcal{M}_{\text{trop}, 0, n}^\alpha(\Delta)$.

7.5 Lemma

Let $n = \#\Delta - 1$. Then $\mathcal{M}_{\text{trop}, 0, n}(\Delta)$ is a polyhedral complex of pure dimension $2n$ and ev is a morphism of polyhedral complexes of the same dimension.

The proof is analogous to the proof of 4.56. We can determine the dimensions of the strata $\mathcal{M}_{\text{trop}, 0, n}^\alpha(\Delta)$ even without using 4.41, because rational curves are regular. (For the proof of 4.41 it was necessary to pass to the relevant subset, and we can therefore not use it here.)

Let us next come to the multiplicity. Note that a 3-valent nonrelevant curve either has a vertex V adjacent to a contracted bounded edge, or a vertex V where two edges point in one direction and the other in the opposite:



In both cases, the multiplicity of V is 0 by definition, and thus so is the multiplicity of the curve. In particular, $\mathcal{M}_{\text{trop}, 0, n}(\Delta) \setminus \widetilde{\mathcal{M}}_{\text{trop}, 0, n}(\Delta)$ consists of curves which do not count anyway. That is, when we count tropical curves, we have

$$N_{\text{trop}}^{\text{irr}}(\Delta, g, \mathcal{P}) = \sum_{C \in \text{ev}^{-1}(\mathcal{P})} \text{mult}(C)$$

also if we now use the evaluation map starting from the bigger space $\mathcal{M}_{\text{trop}, 0, n}(\Delta)$ as defined in 7.4.

In 4.49 we have seen that the multiplicity of a relevant curve is 0 if and only if the evaluation map is not injective restricted to the corresponding stratum $\mathcal{M}_{\text{trop}, 0, n}^\alpha(\Delta)$. This statement is also true for nonrelevant (rational) curves.


7.6 Proposition

Let C be a 3-valent curve of type α in $\mathcal{M}_{\text{trop}, 0, n}(\Delta)$, and let $n = \#\Delta - 1$. Then $\text{mult } C = 0$ if and only if the evaluation map restricted to $\mathcal{M}_{\text{trop}, 0, n}^\alpha(\Delta)$ is not injective.

Proof:

As the statement holds for relevant curves by 4.49, we can assume that $C = (\Gamma, h, x_1, \dots, x_n)$ is not relevant. Assume first that C has a contracted bounded edge. Then both its multiplicity is 0, and the evaluation map is not injective on the corresponding stratum. So the claim holds in this case. Assume next that C has a vertex V as above, where the edges adjacent to it do not span \mathbb{R}^2 . In particular, we have $\text{mult } C = 0$. So we have to see that $\text{ev}|_{\mathcal{M}_{\text{trop}, 0, n}^\alpha(\Delta)}$ is not injective. Note that the image of V and the three edges adjacent to V is just a straight line. Assume all three edges adjacent to V

are unbounded. Then V is the only vertex of the curve. Therefore V has to be the root vertex and the position $h(V)$ is a coordinate of the space $\mathcal{M}_{\text{trop}, 0, n}^\alpha(\Delta)$. But then we can vary this coordinate without changing the image $h(\Gamma)$ of the curve. Assume not all three edges adjacent to V are unbounded. Then we can vary the length of the bounded edge adjacent to V without changing the image $h(\Gamma)$ of the curve. In both cases, we can see that $\text{ev} |_{\mathcal{M}_{\text{trop}, 0, n}^\alpha(\Delta)}$ is not injective.



The converse holds of course, too: for every nonrelevant curve C for which $\text{ev} |_{\mathcal{M}_{\text{trop}, 0, n}^\alpha(\Delta)}$ is not injective we have $\text{mult } C = 0$, already just because it is not relevant. \square

7.7 Remark

Note that for rational curves the sets $\widehat{\mathcal{M}}_{\text{trop}, 0, n}^\alpha(\Delta)$ defined in 4.62 are equal to the sets $\mathcal{M}_{\text{trop}, 0, n}^\alpha(\Delta)$ and the map f_α defined in 4.63 is equal to a matrix representation of the evaluation map. In particular, lemma 4.68 shows in this case that $\text{mult}_{\text{ev}} C = \text{mult } C$ (see remark 4.69). (We required the curves to be relevant in the statement of lemma 4.68, but this requirement was only needed because we applied proposition 4.49 for the case that $\text{mult } C = 0$. But we have seen above in 7.6 that the analogous statement of this proposition holds for rational nonrelevant curves, too.)

That is, for rational curves we have

$$N_{\text{trop}}^{\text{irr}}(\Delta, g, \mathcal{P}) = \sum_{C \in \text{ev}^{-1}(\mathcal{P})} \text{mult}(C) = \sum_{C \in \text{ev}^{-1}(\mathcal{P})} \text{mult}_{\text{ev}} C = \text{deg}_{\text{ev}}(\mathcal{P}).$$

In particular, we have a different formulation of theorem 4.53: for rational curves, it states that the map $\mathcal{P} \mapsto \text{deg}_{\text{ev}}(\mathcal{P})$ is constant. Lemma 4.58, remark 4.59 and lemma 4.72 are enough to prove that $\text{deg}_{\text{ev}}(\mathcal{P})$ does not depend on \mathcal{P} for rational curves.

7.2. TROPICAL FORGETFUL MAPS

Having defined the moduli space we want to work with in this chapter and reconsidered the results of chapter 4 for this space, we can now define tropical forgetful maps.

As in the context of stable maps there are many forgetful maps: for a tropical curve $(\Gamma, h, x_1, \dots, x_n)$ we can forget the map h to \mathbb{R}^2 , or some of the marked points, or both.

7.8 Definition

Let $n' \leq n$ and let $C = (\Gamma, h, x_1, \dots, x_n) \in \mathcal{M}_{\text{trop}, 0, n}(\Delta)$ be an n -marked tropical curve.

Let $C(n')$ be the minimal connected subgraph of Γ that contains the unbounded edges $x_1, \dots, x_{n'}$. Note that $C(n')$ cannot contain vertices of valence 1. So if we “straighten” the graph $C(n')$ at all 2-valent vertices (that is, we replace the two adjacent edges and the vertex by one edge whose length is the sum of the lengths of the original edges) then we obtain an element of $\mathcal{M}_{\text{trop}, 0, n'}$ that we denote by $\text{ft}_{n'}(C)$.

So we can define the *tropical forgetful map* (which forgets some marked points and the map)

$$\text{ft}_{n'} : \mathcal{M}_{\text{trop}, 0, n}(\Delta) \rightarrow \mathcal{M}_{\text{trop}, 0, n'} : C \mapsto \text{ft}_{n'}(C).$$

(Of course, we can also forget any subset of $\{x_1, \dots, x_n\}$ different from $\{x_{n'+1}, \dots, x_n\}$.)

7.9 Definition

Let $n' \leq n$ and let $C = (\Gamma, h, x_1, \dots, x_n) \in \mathcal{M}_{\text{trop}, 0, n}(\Delta)$ be an n -marked tropical curve.

Let $\tilde{C}(n')$ be the minimal connected subgraph of Γ that contains all unmarked ends as well as the marked points $x_1, \dots, x_{n'}$. Again, $\tilde{C}(n')$ cannot have vertices of valence 1. If we straighten $\tilde{C}(n')$ as in 7.8 we obtain an abstract tropical curve $\tilde{\Gamma}$ with $\#\Delta + n'$ unbounded edges. Note that the restriction of h to $\tilde{\Gamma}$ still satisfies the requirements for a parametrized tropical curve, that is $(\tilde{\Gamma}, h|_{\tilde{\Gamma}}, x_1, \dots, x_{n'})$ is an element of $\mathcal{M}_{\text{trop}, 0, n'}(\Delta)$. We denote it by $\tilde{\text{ft}}_{n'}(C)$.


Again, we define the *tropical forgetful map* (which forgets only some marked points)

$$\tilde{\text{ft}}_{n'} : \mathcal{M}_{\text{trop}, 0, n}(\Delta) \rightarrow \mathcal{M}_{\text{trop}, 0, n'}(\Delta) : C \mapsto \tilde{\text{ft}}_{n'}(C).$$


(As before, any other subset of $\{x_1, \dots, x_n\}$ can also be forgotten.)

7.10 Example

The following picture shows a parametrized tropical curve $C = (\Gamma, h, x_1, \dots, x_6)$ of degree 2. The subgraph $C(4)$ is indicated in bold in the graph Γ . Below, the “straightened” version is shown, that is, the image $\text{ft}_4(C) \in \mathcal{M}_{\text{trop}, 0, 4}$.



C



7.11 Lemma

The forgetful maps from definition 7.8 and 7.9 are morphisms of polyhedral complexes.

Proof:

By 7.5 we know that $\mathcal{M}_{\text{trop}, 0, n}(\Delta)$ is a polyhedral complex. The cells are given by the sets $\mathcal{M}_{\text{trop}, 0, n}^\alpha(\Delta)$ depending on the combinatorial type α . For an abstract tropical curve, the combinatorial type is given by the homeomorphism class of the graph Γ . We conclude that the set of all abstract tropical curves of the same combinatorial type α is also a polyhedron, where the coordinates are given by the lengths of the bounded edges. (There is no coordinate of the position of the root vertex here, because we do not have a map to \mathbb{R}^2 but only an abstract tropical curve.) Hence it can be shown analogously that $\mathcal{M}_{\text{trop}, 0, n}$ is a polyhedral complex. Obviously, a forgetful map (no matter of which type — if it forgets points and map or only points) maps a stratum $\mathcal{M}_{\text{trop}, 0, n}^\alpha(\Delta)$ into a cell of curves of the same combinatorial type. The map is furthermore linear, because the coordinates of the target space — the length of the bounded edges of the graph $\text{ft}_{n'}(C)$ (respectively, $\tilde{\text{ft}}_{n'}(C)$) — are given as sums of the coordinates of $\mathcal{M}_{\text{trop}, 0, n}^\alpha(\Delta)$ by construction. \square

In contrast to the classical proof of Kontsevich's formula, we cannot pullback two divisors of $\mathcal{M}_{\text{trop}, 0, 4}$ here, as the concept of tropical divisors is not yet fully developed. Instead, we will combine the evaluation map (which sends two marked points to lines and the other ones to points) with the forgetful map to $\mathcal{M}_{\text{trop}, 0, 4}$. We will show that the degree of this combined evaluation and forgetful map is constant. This statement (see proposition 7.14) replaces the equivalent divisors which we received by pulling back in the classical proof.

The combination of evaluation and forgetful maps we need to consider for Kontsevich's formula is the following:

7.12 Definition

Fix $d \geq 2$, and let $n = 3d$. We set

$$\pi := \text{ev}_1^1 \times \text{ev}_2^2 \times \text{ev}_3 \times \cdots \times \text{ev}_n \times \text{ft}_4 : \mathcal{M}_{\text{trop}, 0, n}(d) \rightarrow \mathbb{R}^{2n-2} \times \mathcal{M}_{\text{trop}, 0, 4},$$

that is π describes the first coordinate of the first marked point, the second coordinate of the second marked point, both coordinates of the other marked points, and the point in $\mathcal{M}_{\text{trop}, 0, 4}$ defined by the first four marked points. Obviously, π is a morphism of polyhedral complexes of pure dimension $2n - 1$.

Note that the first two evaluations — the first coordinate of the first marked point and the second coordinate of the second marked point — determine two lines to which the


first two marked points shall be mapped: the first one parallel to the y -axis, the second one parallel to the x -axis.

7.13 Example

Let us determine the last row of the matrix representing π on the subset of types as in example 7.10 (that is, the row corresponding to the coordinate of $\mathcal{M}_{\text{trop}, 0,4}$). It measures the length of the bounded edge of $\text{ft}_4(C)$. We straightened the two bounded edges l_4 and l_5 to come to the curve $\text{ft}_4(C)$. Hence the last row is

$$(0 \ 0 \ 0 \ 0 \ 0 \ 1 \ 1 \ 0 \ 0 \ 0 \ 0).$$

As a complete example, let us consider the following 4-marked tropical curve with 4 nonmarked ends.



Let us find a matrix representation of π for the stratum of tropical curves of the same type. (We choose a curve of another degree than d here in order to keep the matrix not too big. The construction for π from $\mathcal{M}_{\text{trop}, 0,n}(\Delta)$ for some $\Delta \neq \Delta_d$ (see definition 3.69) is analogous.) The root vertex and the length coordinates are indicated in the picture. We get the following matrix:

$$\begin{pmatrix} 1 & 0 & 0 & -1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 1 \\ 1 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{pmatrix}$$

Note that the determinant of the matrix representation of π has nothing to do with the multiplicity (in contrast to the determinant of a matrix representation of ev) due to the presence of the last row corresponding to $\mathcal{M}_{\text{trop}, 0,4}$ and because we evaluate the first two points at lines only.

Now we come to the central result of this section, the statement that the degrees $\text{deg}_\pi(\mathcal{P})$ (see definition 4.55) of π do not depend on \mathcal{P} . Here, \mathcal{P} denotes an element of the space


$\mathbb{R}^{2n-2} \times \mathcal{M}_{\text{trop}, 0,4}$. That is, $\mathcal{P} = (a, b, p_3, \dots, p_n, C')$ where a denotes the x -coordinate of a line parallel to the y -axis, b the y -coordinate of a line parallel to the x -axis, p_3, \dots, p_n denote $3d - 2$ points and C' an element in $\mathcal{M}_{\text{trop}, 0,4}$.

7.14 Proposition

The degrees $\text{deg}_\pi(\mathcal{P})$ of the map π defined in 7.12 do not depend on \mathcal{P} (as long as \mathcal{P} is in π -general position).

Proof:

The proof is analogous to the proof of theorem 4.53 which shows that the number $N_{\text{trop}}^{\text{irr}}(\Delta, g, \mathcal{P})$ (which is here in the case of rational curves equal to $\text{deg}_{\text{ev}}(\mathcal{P})$, see remark 7.7) does not depend on \mathcal{P} . An essential ingredient for this proof was lemma 4.72. The analogous idea will be used here, too. Analogously to 4.58, $\text{deg}_\pi(\mathcal{P})$ is locally constant on the set of points in π -general position. So with the same arguments as in remark 4.59, we only have to consider a general point of the subset of points which are not in π -general position. Such a point is the image under π of a curve C which is not in general position in $\mathcal{M}_{\text{trop}, 0,n}(d)$, that is, a curve of codimension 1. As we are working with rational curves, the only type of codimension 1 is a type with exactly one 4-valent vertex. Let C be such a curve and let $\mathcal{P}' = \pi(C)$. Analogously to 4.72 we know the three types α_1, α_2 and α_3 which have C in their boundary:



To check that $\text{deg}_\pi(\mathcal{P})$ is constant, we only have to see that the sum of the π -multiplicities of the inverse images of a point configuration \mathcal{P}'' near \mathcal{P}' in the strata $\mathcal{M}_{\text{trop}, 0,n}^{\alpha_i}(d)$, $i = 1, 2, 3$, does not depend on \mathcal{P}'' . Let A_i be a matrix representation of $\pi|_{\mathcal{M}_{\text{trop}, 0,n}^{\alpha_i}(d)}$.

As in the proof of 4.72, the first step is to prove that the sum $\det A_1 + \det A_2 + \det A_3 = 0$ and the second step is to show which types can occur in a preimage.

The following table represents all three matrices A_1, A_2 and A_3 .


	$h(V)$	l_1	l_2	l_3	l_4	l^{α_1}	l^{α_2}	l^{α_3}
points behind e_1	E_2	v_1	0	0	0	0	0	0
points behind e_2	E_2	0	v_2	0	0	0	$v_2 + v_3$	$v_2 + v_4$
points behind e_3	E_2	0	0	v_3	0	$v_3 + v_4$	$v_2 + v_3$	0
points behind e_4	E_2	0	0	0	v_4	$v_3 + v_4$	0	$v_2 + v_4$
coordinate of $\mathcal{M}_{\text{trop}, 0,4}$	0	*	*	*	*	**	**	**

Each matrix contains the first block of columns, and the i -th of the last three columns. The first columns represent the coordinate $h(V)$ of the root vertex, the following four the coordinates of the lengths of the edges e_1, \dots, e_4 . The last three correspond to the length

of the new edge e in the three different types. The last row corresponds to the coordinate in $\mathcal{M}_{\text{trop}, 0,4}$. Each $*$ and $**$ stands for 0 or 1 (which will be explained further below).


To look at the entries marked $*$ further we will distinguish several cases depending on how many of the edges e_1, \dots, e_4 of C are contained in the subgraph $C(4)$ of definition 7.8:

- (1) 4 edges: Then $\text{ft}_4(C)$ — the last coordinate of \mathcal{P}' — is the curve (4) of example 4.9 (that is, all four unbounded edges x_1, \dots, x_4 come together at a vertex), and the three types $\alpha_1, \alpha_2, \alpha_3$ are mapped precisely to the three other types (1), (2), (3) of $\mathcal{M}_{\text{trop}, 0,4}$ by ft_4 (that differ by which of the four unbounded edges x_1, \dots, x_4 come together at the two 3-valent vertices — see example 4.9). Hence the three types $\alpha_1, \alpha_2, \alpha_3$ are mapped to the three cells of $\mathbb{R}^{2n-2} \times \mathcal{M}_{\text{trop}, 0,4}$ around \mathcal{P}' by π . For these three types the length parameter in $\mathcal{M}_{\text{trop}, 0,4}$ is simply the length of the newly inserted edge e . Hence the entries $*$ in the matrix above are all 0, whereas the entries $**$ are all 1. It follows that the three matrices A_1, A_2, A_3 have in the last row a one as the bottom right entry and only zeroes else. Therefore their determinants do not depend on the last column. But this is the only column that differs for the three matrices. Hence A_1, A_2 and A_3 all have the same determinant. The last coordinate (the $\mathcal{M}_{\text{trop}, 0,4}$ -coordinate) C'' of a point configuration \mathcal{P}'' near \mathcal{P}' is a curve of precisely one of the types (1), (2) or (3) of 4.9. Hence also precisely one of the types $\alpha_1, \alpha_2, \alpha_3$ occurs in the preimage under π of \mathcal{P}'' . As the π -multiplicity of all three types is equal, it follows that deg_π is locally constant around C . This completes the proof of the proposition in this case.
- (2) 3 edges: The following picture shows what the combinatorial types $\alpha, \alpha_1, \alpha_2, \alpha_3$ look like locally around the vertex V in this case. As in example 7.10 we have drawn the edges belonging to $C(4)$ bold.



We can see that exactly one edge e_i (e_4 in the picture above) counts towards the length parameter in $\mathcal{M}_{\text{trop}, 0,4}$, and that the newly inserted edge counts towards this length parameter in exactly one of the combinatorial types α_i (α_1 in the picture above). Hence in the table showing the matrices A_i , exactly one of the entries $*$ and exactly one of the entries $**$ is 1, whereas the others are 0.

- (3) 2 edges: Assume first that 2 marked points lie behind one e_i and 2 behind another.



Then we can see that two edges (e_1 and e_4 in the picture) contribute to the length parameter in $\mathcal{M}_{\text{trop}, 0, 4}$, and the inserted edge contributes in two of the combinatorial types (α_1 and α_3). That is, 2 of the entries $*$ and also 2 of the entries $**$ above are 1, whereas the others are 0. If there are three marked points behind one of the e_i and one behind another, then none of the edges contributes, and all $*$ and $**$ entries are 0.

- (4) 1 edge: It is not possible that exactly one of the edges e_i is contained in $C(4)$.
- (5) 0 edges: Then neither the edges e_1, \dots, e_4 nor the newly inserted edge contribute to the length parameter in $\mathcal{M}_{\text{trop}, 0, 4}$. That is, all the $*$ and $**$ entries are 0.

To sum up, we have seen in any case that there are equally many entries $**$ equal to 1 as there are entries $*$ equal to 1. So using the same operations as in 4.72 — adding the last three columns to get a matrix whose determinant is equal to $\det A_1 + \det A_2 + \det A_3$, subtracting the four l_i -columns and v_1 times the $h(V)$ -columns — we can see that $\det A_1 + \det A_2 + \det A_3$ is equal to the determinant of a matrix with a 0-column, hence is 0. Using the same analysis as in 4.72 we can also see that the question whether there is a point in $\pi|_{\mathcal{M}_{\text{trop}, 0, n}^{\alpha_i}(d)}^{-1}(\mathcal{P}'')$ or not depends only on the sign of $\det A_i$. Hence $\mathcal{P} \mapsto \deg_{\pi}(\mathcal{P})$ is locally constant at \mathcal{P}' . \square

7.3. REDUCIBLE CURVES AND KONTSEVICH'S FORMULA

We have seen in 7.14 that the degrees of the morphism π defined in 7.12 do not depend on the chosen point configuration on the target $\mathbb{R}^{2n-2} \times \mathcal{M}_{\text{trop}, 0, 4}$. Now we want to apply this result by taking two different point configurations \mathcal{P}_1 and \mathcal{P}_2 and equating the two degrees $\deg_{\pi}(\mathcal{P}_1) = \deg_{\pi}(\mathcal{P}_2)$. We want to choose \mathcal{P}_1 and \mathcal{P}_2 such that the $\mathcal{M}_{\text{trop}, 0, 4}$ -coordinate is very large, but corresponds to curves of different types, namely to (1) and (2) in example 4.9. That is, if we consider the abstract tropical curve $\text{ft}_4(C)$ for a curve C in $\pi^{-1}(\mathcal{P}_1)$, then the two ends x_1 and x_2 come together at a 3-valent vertex, whereas for a curve in $\pi^{-1}(\mathcal{P}_2)$, the ends x_1 and x_3 come together.

First, we need to see that a very large $\mathcal{M}_{\text{trop}, 0, 4}$ -coordinate implies that the curves in $\pi^{-1}(\mathcal{P})$ contain a contracted bounded edge, that is, an edge which is mapped to a point by h .

7.15 Proposition

Let $d \geq 2$ and $n = 3d$, and let $\mathcal{P} \in \mathbb{R}^{2n-2} \times \mathcal{M}_{\text{trop}, 0, 4}$ be a point in π -general position whose $\mathcal{M}_{\text{trop}, 0, 4}$ -coordinate is very large (so that it corresponds to a 4-marked curve of type (1), (2), or (3) as in example 4.9 with a very large length l of the bounded edge).

Then every tropical curve $C \in \pi^{-1}(\mathcal{P})$ with $\text{mult}_{\pi}(C) \neq 0$ has a contracted bounded edge. This contracted edge is contained in the subgraph $C(4)$.

Proof:

We have to show that the set of all points $\text{ft}_4(C) \in \mathcal{M}_{\text{trop}, 0, 4}$ is bounded in $\mathcal{M}_{\text{trop}, 0, 4}$, where C runs over all curves in $\mathcal{M}_{\text{trop}, 0, n}(d)$ with non-zero π -multiplicity that have no contracted bounded edge and where the marked points are mapped as required to the two


lines and $3d-2$ points of \mathcal{P} . As there are only finitely many combinatorial types by lemma 4.30 we can restrict our considerations to curves of a fixed (but arbitrary) combinatorial type α . Since \mathcal{P} is in π -general position we can assume that the codimension of α is 0, that is C is 3-valent.

Let $C_1 \in \mathcal{M}_{\text{trop}, 0, n-2}(d)$ be the curve that is obtained from C by forgetting the first two marked points as in definition 7.9. We claim that C_1 has exactly one string (see definition 4.46). (As C_1 is rational, this string will then in fact be a path connecting two unbounded edges.) C_1 must have at least one string by remark 4.50, because C_1 has less than $3d-1 = n-1$ marked points. So assume C_1 has two or more strings. Then C_1 would move in an at least 2-dimensional family with the images of x_3, \dots, x_n fixed. Hence also C moves in an at least 2-dimensional family with the x -coordinate of $h(x_1)$, the y -coordinate of $h(x_2)$, and both coordinates of $h(x_3), \dots, h(x_n)$ fixed. As $\mathcal{M}_{\text{trop}, 0, 4}$ is one-dimensional this means that C moves in an at least one-dimensional family with the image point under π fixed. Hence π is not a local isomorphism and so $\text{mult}_\pi(C) = 0$ which contradicts our assumption.

So let Γ_1 be the unique string in C_1 . The deformations of C_1 where the images of the marked points are fixed are then precisely the ones of the string described in the proof of proposition 4.49. The edges adjacent to Γ_1 must be bounded since otherwise there would be two strings. So if there are edges adjacent to Γ_1 to both sides of Γ_1 as in picture (a) below (note that there are no contracted bounded edges by assumption) then the “movement” of the string is bounded. (This is true because if we move the string to either side, we can only move until the length of one of the adjacent bounded edges shrinks to 0, leading to a different combinatorial type.) That is, the deformations of C_1 with fixed combinatorial type and fixed images of the marked points are bounded on both sides. Now let us consider the deformations of C again, where we fix the combinatorial type and ask the marked points to map to the first coordinates (a, b, p_3, \dots, p_n) of \mathcal{P} under π — that is, the first two to the two lines which are fixed by \mathcal{P} , and the others to the fixed points. For these deformations of C we have that the lengths of all inner edges are bounded except possibly the edges adjacent to x_1 and x_2 . This is true, because the bounded edges adjacent to x_1 and x_2 may have become unbounded when forgetting x_1 and x_2 . However, this only happens if one of the edges adjacent to x_1 is unbounded, too. No two of the marked points x_1, \dots, x_4 can be adjacent to the same vertex, because otherwise C would have a contracted edge — the third edge adjacent to the same vertex would be of direction 0 due to the balancing condition. Hence if x_1 is adjacent to another unbounded edge, then this is not one of the marked points. But then the length of the bounded edges adjacent to x_1 and x_2 cannot contribute to the length parameter. As all other lengths are bounded, the image of these deformations of C under ft_4 is bounded in $\mathcal{M}_{\text{trop}, 0, 4}$ as well.

So we only have to consider the case when all adjacent edges of Γ_1 are on the same side of Γ_1 , say on the right side as in picture (b) below. Label the edges of Γ_1 (respectively, their direction vectors) by v_1, \dots, v_k and the adjacent edges of the curve by w_1, \dots, w_{k-1} as in the picture. As above the movement of the string of C_1 to the right within its combinatorial type is bounded. If one of the directions w_{i+1} is obtained from w_i by a left turn (as it is the case for $i = 1$ in the picture) then the edges w_i and w_{i+1} meet on the left

of Γ_1 . This restricts the movement of the string of C_1 to the left within its combinatorial type, too, since the corresponding edge v_{i+1} then shrinks to length 0. Therefore we have again as in case (a) above that the image of these curves under ft_4 in $\mathcal{M}_{\text{trop}, 0,4}$ is bounded.



So we can assume that for all i the direction w_{i+1} is either the same as w_i or obtained from w_i by a right turn as in picture (c). The balancing condition then shows that for all i both the directions v_{i+1} and $-w_{i+1}$ lie in the angle between v_i and $-w_i$ (shaded in the picture above). Therefore, all directions v_i and $-w_i$ lie within the angle between v_1 and $-w_1$. In particular, the image of the string Γ_1 cannot have any self-intersections in \mathbb{R}^2 . We can therefore pass to the (local) dual picture (d) (see 2.15) where the edges dual to w_i correspond to a concave side of the polygon whose other two edges are the ones dual to v_1 and v_k .

But note that there are no such concave polygons *with integer vertices* if the two outer edges (dual to v_1 and v_k) are two of the vectors $\pm(1, 0)$, $\pm(0, 1)$, $\pm(1, -1)$ that can occur as dual edges of an unbounded edge of a tropical curve of degree d . Therefore the string can consist at most of the two unbounded ends v_1 and v_2 that are connected to the rest of the curve by exactly one bounded edge w_1 . This situation is shown in picture (e).

In this case the movement of the string is indeed not bounded to the left. Note that then w_1 is the only internal edge whose length is not bounded within the deformations of C_1 . But this unbounded length of the edge w_1 cannot count towards the length parameter in $\mathcal{M}_{\text{trop}, 0,4}$ for the deformations of C as this would require two of the marked points x_1, \dots, x_4 to lie on v_1 or v_2 for all curves in the deformation, which is incompatible with the lines and points given by \mathcal{P} that C is required to meet. Therefore the image of these curves under ft_4 is bounded in $\mathcal{M}_{\text{trop}, 0,4}$, too.


As the set of points $\text{ft}_4(C)$ in $\mathcal{M}_{\text{trop}, 0,4}$ is bounded, where C goes over all curves without contracted bounded edges, we can conclude that for a very large length parameter in $\mathcal{M}_{\text{trop}, 0,4}$, the preimages have to contain a contracted bounded edge e . The edge e must be contained in the subgraph $C(4)$ that we straightened to get $\text{ft}_4(C) = C'$ (see definition 7.8), because otherwise we could not have a very large length parameter in $\mathcal{M}_{\text{trop}, 0,4}$.

□

So we know that for the two configurations \mathcal{P}_1 and \mathcal{P}_2 , for which we want to compare the degrees of π , the inverse images under π are curves which contain a contracted bounded edge. Let us describe what a curve with a contracted bounded edge looks like.

7.16 Remark

Let $C = (\Gamma, h, x_1, \dots, x_n)$ be a tropical curve with a contracted bounded edge e , and assume that there is at least one more bounded edge on both sides of e . Then we can cut Γ at e in the same way as in the proof of lemma 4.68 — we replace e on both sides by an unbounded edge which is then also contracted to a point. Let us consider these two new contracted unbounded edges as new marked points x_{n+1} and x_{n+2} . Restricting h to the two graphs Γ_A and Γ_B we obtain in this way, we get two tropical curves C_A and C_B . They are then curves of degree d_A respectively d_B , where $d_A + d_B = d$ and $d_A, d_B < d$. This is true due to the balancing condition. We will say in this situation that C is obtained by *gluing* C_A and C_B along the identification $x_{n+1} = x_{n+2}$, and that C is a *reducible* tropical curve that can be *decomposed* into C_A and C_B (even though the graph Γ is not really disconnected here, it can only be made disconnected). For the image in \mathbb{R}^2 we obviously have $h(\Gamma) = h(\Gamma_A) \cup h(\Gamma_B)$. That is, when we consider the image $h(\Gamma)$ as an unparametrized tropical curve, then it is indeed reducible (see 5.6). This is true because the unique simple parametrization will consist of the two connected components Γ_A and Γ_B , and will not contain the contracted bounded edge e . That is, C considered as unparametrized tropical curve is in fact just the union of the two curves C_A and C_B of smaller degree.



By 7.15 we know that each curve in the preimage $\pi^{-1}(\mathcal{P}_1)$ respectively $\pi^{-1}(\mathcal{P}_2)$ contains an edge which is contracted to a point. That is, by remark 7.16 there will be reducible curves in such a preimage. Let us describe the preimage of \mathcal{P}_1 more detailed. The situation for \mathcal{P}_2 is analogous.

7.17 Lemma

Let $\mathcal{P}_1 = (a, b, p_3, \dots, p_n, C') \in \mathbb{R}^{2n-2} \times \mathcal{M}_{\text{trop}, 0, 4}$ be a point in π -general position such that $C' \in \mathcal{M}_{\text{trop}, 0, 4}$ is of type (1) (that is, in C' the two marked ends x_1 and x_2 come together at a 3-valent vertex — see example 4.9) with a very large length parameter.

Then for every tropical curve C in $\pi^{-1}(\mathcal{P}_1)$ with non-zero π -multiplicity we have exactly one of the following cases:

- (1) x_1 and x_2 are adjacent to the same vertex (that maps to (a, b) under h);
- (2) C is reducible and decomposes uniquely into two components C_A and C_B (as in 7.16) of some degrees d_A and d_B with $d_A + d_B = d$ such that the marked points x_1 and x_2 are on C_A , the points x_3 and x_4 are on C_B , and exactly $3d_A - 1$ of the other points x_5, \dots, x_n are on C_A .

Proof:

By 7.15 we know that any curve $C \in \pi^{-1}(\mathcal{P}_1)$ with non-zero π -multiplicity has at least one contracted bounded edge. In fact, C cannot have more edges which are contracted: if C had at least two contracted edges, then there would be $2n - 2$ coordinates in the target of π (the evaluation maps) that depend only on $2n - 3$ variables (namely the root vertex and the lengths of all but the two contracted of the $2n - 3$ bounded edges), hence we would have $\text{mult}_\pi(C) = 0$.

So let e be the unique contracted bounded edge of C . Recall that e must be contained in the subgraph $C(4)$ that we straightened to get $\text{ft}_4(C) = C'$ (see definition 7.8), because else we could not have a very large length parameter in $\mathcal{M}_{\text{trop}, 0, 4}$. As the point C' is of type (1) the two ends x_1 and x_2 must be to one side, and x_3 and x_4 to the other side of e . Denote these sides by C_A and C_B , respectively.

Assume first that there are no bounded edges in C_A . Then C is not reducible as in remark 7.16. Instead C_A consists only of e , x_1 and x_2 . This means we are in case 1. The incidence conditions then require that all of C_A must be mapped to the intersection point (a, b) of the two lines $\{x = a\}$ and $\{y = b\}$ that were prescribed by \mathcal{P}_1 . (Note that it is not possible that there are no bounded edges in C_B , because else x_3 and x_4 would be mapped to the same point in \mathbb{R}^2 .)

Now assume that there are bounded edges to both sides C_A and C_B of e . In this case C is reducible as in remark 7.16, so we are in case 2. First we claim that the two marked points x_1 and x_2 are not adjacent to the same vertex V . If they were, the third edge adjacent to V would be mapped to a point due to the balancing condition, which contradicts the above where we have seen that e is the only contracted edge. Now let n_1 (and n_2) be the number of marked points x_5, \dots, x_n on C_A (respectively C_B). We have to show that $n_1 = 3d_A - 1$ and $n_2 = 3d_B - 3$. So assume that $n_1 \geq 3d_A$. Then at least $2n_1 + 2 \geq 3d_A + n_1 + 2$ of the coordinates of π (the images of the n_1 marked points, the first image coordinate of x_1 and the second of x_2) would depend only on $3d_A + n_1 + 1$ coordinates (2 for the root vertex and one for each of the $3d_A + (n_1 + 2) - 3$ for the bounded edges which are part of C_A), leading to a zero π -multiplicity. Hence we have $n_1 \leq 3d_A - 1$. Analogously, we can see that $n_2 \leq 3d_B - 3$. As the total number of marked points is $n_1 + n_2 = n - 4 = (3d_A - 1) + (3d_B - 3)$, we must therefore have $n_1 = 3d_A - 1$ and $n_2 = 3d_B - 3$. \square

Lemma 7.17 tells us that we can interpret the inverse images $\pi^{-1}(\mathcal{P}_1)$ — at least under some assumptions — as reducible curves, just as we can interpret the stable maps in the divisor $D_{1,2/3,4} = \sum_{d_A+d_B=d} \sum_{A,B} D(d_A, d_B, A, B)$ defined in 3.28 as reducible curves. The idea of the “tropical” proof of Kontsevich’s formula is indeed quite analogous to the classical situation described in section 3.2. As in the classical situation, we would like to count the components of the reducible curves in $\pi^{-1}(\mathcal{P}_1)$ separately, and the number of ways two such components can be attached to each other. Before we can do this, we have to study whether each choice of two components C_A and C_B and a way to attach them to each other really yields a curve C in $\pi^{-1}(\mathcal{P}_1)$. This is shown in the following remark:

7.18 Remark

As in lemma 7.17 let $\mathcal{P}_1 = (a, b, p_3, \dots, p_n, C') \in \mathbb{R}^{2n-2} \times \mathcal{M}_{\text{trop}, 0,4}$ be a point in π -general position such that $C' \in \mathcal{M}_{\text{trop}, 0,4}$ is of type (1) (that is, in C' the two marked ends x_1 and x_2 come together at a 3-valent vertex — see example 4.9) with a very large length parameter. Let C_A and C_B be two unmarked tropical curves of degree d_A respectively d_B with $d_A + d_B = d$ such that the image of C_A passes through the lines $L_1 := \{x = a\}$, $L_2 := \{y = b\}$ and $3d_A - 1$ of the points p_5, \dots, p_n , and the image of C_B through p_3, p_4 , and the remaining $3d_B - 3$ of the points p_5, \dots, p_n .

Then for each choice of points $x_{n+1} \in C_A$ and $x_{n+2} \in C_B$ that map to the same image point in \mathbb{R}^2 , and for each choice of points x_1, \dots, x_n on C_A and C_B that map to $L_1, L_2, p_3, \dots, p_n$, respectively, we can make C_A and C_B into marked tropical curves by adding marked unbounded edges at the points x_i . Then, we can replace the two marked unbounded edges x_{n+1} and x_{n+2} by one contracted bounded edge e and glue in that way the two curves C_A and C_B together to a reducible n -marked curve C in $\pi^{-1}(\mathcal{P}_1)$ — “converse” to what we have done in remark 7.16. (The length of the contracted edge e is determined by C' , hence we really get one single reducible curve C in that way.)

As \mathcal{P}_1 is required to be in π -general position we can never construct a curve C in this way which is not 3-valent. In particular, C_A and C_B are 3-valent. Moreover, a point that is in the image of both C_A and C_B cannot be a vertex of one of the curves. Hence it is not possible that C_A and C_B share a common line segment in \mathbb{R}^2 . In the same way we can see that the image of C_A cannot meet L_1 or L_2 in a vertex or have a line segment in common with L_1 or L_2 , and cannot meet $L_1 \cap L_2$ at all.

To sum up, we have seen that after choosing the two curves C_A and C_B as well as the points x_1, \dots, x_n and x_{n+1}, x_{n+2} on them, there is a unique curve C in $\pi^{-1}(\mathcal{P}_1)$ obtained from this data. In order to compute the degree of π at \mathcal{P}_1 , we have to sum over all points in $\pi^{-1}(\mathcal{P}_1)$. Instead, for the curves of type 2 in lemma 7.17 we can as well sum over all choices of $C_A, C_B, x_1, \dots, x_n, x_{n+1}, x_{n+2}$ as above.

Before we can really count the curves $C \in \pi^{-1}(\mathcal{P}_1)$ (respectively, the choices of $C_A, C_B, x_1, \dots, x_n, x_{n+1}, x_{n+2}$ as in remark 7.18) we have to compute the multiplicity of π at the curve C , and compare it with the multiplicities of C_A and C_B which we would like to count instead. As already mentioned in remark 7.18, we want to sum over all choices of $C_A, C_B, x_1, \dots, x_n, x_{n+1}, x_{n+2}$, where $x_{n+1} = x_{n+2}$ is a common point of C_A and C_B , and (for $i = 1, 2$) x_i is a common point of C_A and L_i . Analogously to the classical situation, an intersection point of C_A and C_B (respectively, of C_A and L_i) can have a multiplicity which we need to consider when we want to count the possibilities to choose $x_{n+1} = x_{n+2}$ respectively x_i . Therefore, we introduce the notion of intersection multiplicity and cite the tropical Bézout’s theorem, before we compare the multiplicities of C, C_A and C_B in proposition 7.21.

7.19 Definition

Let C_A and C_B be two (parametrized) tropical curves and let $p \in C_A \cap C_B$ be a common point of the images of both curves in \mathbb{R}^2 . Assume $\#h_A^{-1}(p) = \#h_B^{-1}(p) = 1$ and both preimages are distinct from the vertices of C_A respectively C_B . We define the *intersection*

multiplicity of C_A and C_B at the point p to be

$$(C_A \cdot C_B)_p = |\det(v, w)|$$

where v and w are the direction vectors of C_A and C_B at p . (The directions at p are well-defined (up to sign), because for both curves, $h^{-1}(p)$ lies on a single edge (of a well-defined direction) by assumption.)

For a line L_i and a marked point x_i ($i = 1, 2$) of C_A (which is adjacent to a 3-valent vertex and) which is mapped to L_i we define $(C_A \cdot L_i)_{x_i}$ to be the first respectively second coordinate of the direction vector of C_A at x_i . (Again, this direction vector is well-defined (up to sign), because the two edges which are adjacent to the same vertex as x_i are mapped to the same line due to the balancing condition.)


7.20 Theorem (Bézout’s theorem)

Let C_A and C_B be two (parametrized) tropical curves of degree d_A and d_B , such that the intersection $C_A \cap C_B$ is finite, for each $p \in C_A \cap C_B$ we have $\#h_A^{-1}(p) = \#h_B^{-1}(p) = 1$, and the preimages are distinct from the vertices of C_A respectively C_B .

Then the number of intersection points of C_A and C_B , counted with the intersection multiplicity as defined in 7.19, is equal to the product of the degrees,

$$\sum_{p \in C_A \cap C_B} (C_A \cdot C_B)_p = d_A \cdot d_B.$$

For a proof, see [25] theorem 4.2. The idea of the proof is to consider first a special position for the two tropical curves C_A and C_B : Assume that the only intersection points of C_A and C_B lie on the unbounded edges of C_A of direction $(-1, 0)$ and the unbounded edges of C_B of direction $(0, -1)$.



It is easy to see that each intersection point is of intersection multiplicity 1. Second, it can be shown that the sum $\sum_{p \in C_A \cap C_B} (C_A \cdot C_B)_p$ stays constant when we move C_A . This is true because the balancing condition is fulfilled at each vertex.

Now we are ready to compare the multiplicity of the curves $C \in \pi^{-1}(\mathcal{P}_1)$ with the multiplicities of the components C_A and C_B :

7.21 Proposition

With the notations as in lemma 7.17 and remark 7.18, let C be a point in $\pi^{-1}(\mathcal{P}_1)$. Then

- (1) if C is of type 1 as in lemma 7.17 its π -multiplicity is $\text{mult}_{\text{ev}}(C_A)$, where C_A denotes the curve obtained from C by forgetting x_1 , and ev is the evaluation at the $3d - 1$ points x_2, \dots, x_n ;
- (2) if C is of type 2 as in lemma 7.17 its π -multiplicity is

$$\text{mult}_{\pi}(C) = \text{mult}_{\text{ev}}(C_A) \cdot \text{mult}_{\text{ev}}(C_B) \cdot (C_A \cdot C_B)_{x_{n+1}=x_{n+2}} \cdot (C_A \cdot L_1)_{x_1} \cdot (C_A \cdot L_2)_{x_2},$$

where $\text{mult}_{\text{ev}}(C_A)$ (respectively $\text{mult}_{\text{ev}}(C_B)$) denotes the multiplicities of the evaluation map at the $3d_i - 1$ points of x_3, \dots, x_n that lie on the respective curve (for the definition of the intersection multiplicities, see 7.19).


Proof:

We know that any curve $C \in \pi^{-1}(\mathcal{P}_1)$ (which counts with a nonzero π -multiplicity) has to be of type 1 respectively 2 as in 7.17. In particular, C contains exactly one contracted bounded edge e which is part of the subgraph $C(4)$ we straightend in order to get $\text{ft}_4(C) = C'$ as in 7.8. In order to prove the statement about the multiplicities, we have to set up the matrix for π and compute its determinant. First note that in both cases 1 and 2 the length of the edge e does not contribute to the evaluation of any marked point, because it is contracted. It contributes with a factor of 1 to the $\mathcal{M}_{\text{trop}, 0, 4}$ -coordinate of π , because it is part of $C(4)$. Hence in the matrix for π , the column corresponding to the contracted bounded edge e has only one entry 1 and all others zero. To compute the determinant of this matrix we may therefore drop both the $\mathcal{M}_{\text{trop}, 0, 4}$ -row and the column corresponding to e .

Let us first consider case 1. Then, in the curve C , the marked points x_1 and x_2 are both adjacent to e and mapped to the same point (a, b) by h . The matrix that we obtain when we delete the $\mathcal{M}_{\text{trop}, 0, 4}$ -row and the e -column is exactly the same as if we had only one marked point instead of x_1 and x_2 and evaluate this point for both coordinates in \mathbb{R}^2 (instead of evaluating x_1 for the first and x_2 for the second). This proves case 1.

In case 2 we know that C is reducible with components C_A and C_B by 7.17. Let us first consider the marked point x_1 , where we only evaluate the first coordinate. Let e_1 and e_2 be the two adjacent edges and assume first that both of them are bounded.

As indicated in the picture below, denote the direction vector of e_1 by $-v = (-v^1, -v^2)$. Due to the balancing condition, the direction of e_2 is v . Denote the lengths of the two bounded edges e_1 and e_2 by l_1 and l_2 . Assume that the root vertex is on the e_1 -side of x_1 — it is also shown in the picture.




Then the entries of the matrix for π corresponding to l_1 and l_2 are as shown in the following table:

↓ evaluation at...	l_1	l_2
x_1 (1 row)	v^1	0
points reached via e_1 from x_1 (2 rows each, except only 1 for x_2)	0	0
points reached via e_2 from x_1 (2 rows each, except only 1 for x_2)	v	v

We see that after subtracting the l_2 -column from the l_1 -column we again get one column with only one non-zero entry v^1 . So for the determinant we get $v^1 = (C_A \cdot L_1)_{x_1}$ as a factor, dropping the corresponding row and column (which simply means forgetting the point x_1 as in definition 7.9). Essentially the same argument holds if one of the adjacent edges — say e_2 — is unbounded: in this case there is only an l_1 -column which has zeroes everywhere except in the one x_1 -row where the entry is v^1 .

The same is of course true for x_2 and leads to a factor of $(C_A \cdot L_2)_{x_2}$.

Next we consider the contracted bounded edge e at which we split the curve C into the two parts C_A and C_B . Choose one of its boundary points as root vertex V . Denote the adjacent edges and their directions as in the following picture:



First assume that all edges e_1, \dots, e_4 are bounded. Let $l_i = l(e_i)$. We assume as before that there are n_1 of the marked points x_5, \dots, x_n on C_A and n_2 on C_B . (Recall that we have forgotten the two marked points x_1 and x_2 already, and cancelled the corresponding columns and rows of the matrix of π , leaving a matrix of size $2n - 4$ altogether. We also cancelled the e -column and the $\mathcal{M}_{\text{trop}, 0,4}$ -row.)

Then the remaining matrix (of size $2n - 4$) reads

		lengths in C_A				lengths in C_B		
		root	$(2n_1 - 3 \text{ cols})$			$(2n_2 + 1 \text{ cols})$		
		E_2	l_1	l_2	l_3	l_4		
$(2n_1$ rows)	points behind e_1	E_2	*	v	0	0	0	0
	points behind e_2	E_2	*	0	$-v$	0	0	0
$(2n_2 + 4$ rows)	points behind e_3	E_2	0	0	0	w	0	*
	points behind e_4	E_2	0	0	0	0	$-w$	*

where E_2 is the 2×2 unit matrix and $*$ denotes arbitrary entries. We know already that $n_1 = 3d_A - 1$ and $n_2 = 3d_B - 3$ due to lemma 7.17. That is, we have $n_1 + 1$ marked points on C_A (counting x_{n+1} , too) and $n_2 + 3$ on C_B (counting x_3, x_4 and x_{n+2} as well). Hence there are $(n_1 + 1 + 3d_A) - 3 = 2n_1 - 1$ bounded edges on C_A and $(n_2 + 3 + 3d_B) - 3 = 2n_2 + 3$ bounded edges on C_B . In the table, we show the columns for l_1, \dots, l_4 separately, so there are $2n_1 - 3$ other columns for the bounded edges of C_A and $2n_2 + 1$ other columns for the bounded edges of C_B .

Add v times the root columns to the l_2 -column, subtract the l_1 -column from the l_2 -column and the l_4 -column from the l_3 -column to obtain the following matrix with the same determinant:

		lengths in C_A			lengths in C_B			
		root	$(2n_1 - 3 \text{ cols})$	l_1	l_2	l_3	l_4	$(2n_2 + 1 \text{ cols})$
$(2n_1$	points behind e_1	E_2	*	v	0	0	0	0
rows)	points behind e_2	E_2	*	0	0	0	0	0
$(2n_2 + 4$	points behind e_3	E_2	0	0	v	w	0	*
rows)	points behind e_4	E_2	0	0	v	w	$-w$	*

Note that this matrix has a block form with a zero block at the top right. Denote the top left block (of size $2n_1$) by A_1 and the bottom right (of size $2n_2 + 4$) by A_2 , so that the multiplicity that we are looking for is $|\det A_1 \cdot \det A_2|$.

The matrix A_1 is precisely the matrix for the evaluation map of C_A if we forget the marked point x_{n+1} that replaces e after cutting. Note that forgetting this marked point means we have to straighten the two edges e_1 and e_2 to get one edge e'_1 . The column in A_1 corresponding to e'_1 is the column still denoted with l_1 . (Of course, V can no longer be the root vertex, as we straightened the graph at V . We have to take the other end point of e_2 instead, which is then also a vertex of the straightened graph.) Hence $|\det A_1| = \text{mult}_{\text{ev}}(C_A)$. In the same way the matrix for the evaluation map of C_B is the matrix A'_2 obtained from A_2 by replacing v and w in the first two columns by the first and second unit vector, respectively. (As before, we have to forget the unbounded edge x_{n+2} which replaces the cut edge e , straighten the graph to make e_3 and e_4 to one edge with column corresponding to l_4 and choose the root vertex to be the other end point of e_3 .) The matrix A_2 is obtained from A'_2 by right multiplication with the matrix

$$\begin{pmatrix} v & w & 0 \\ 0 & 0 & E_{2n_2-4} \end{pmatrix}$$

which has determinant $\det(v, w)$. So we conclude that

$$|\det A_2| = |\det(v, w)| \cdot |\det A'_2| = (C_A \cdot C_B)_{x_{n+1}=x_{n+2}} \cdot \text{mult}_{\text{ev}}(C_B).$$

Altogether we have

$$\begin{aligned} \text{mult}_\pi(C) &= (C_A \cdot L_1)_{x_1} \cdot (C_A \cdot L_2)_{x_2} \cdot |\det A_1| \cdot |\det A_2| \\ &= (C_A \cdot L_1)_{x_1} \cdot (C_A \cdot L_2)_{x_2} \cdot \text{mult}_{\text{ev}}(C_A) \cdot |\det(v, w)| \cdot |\det A'_2| \\ &= (C_A \cdot L_1)_{x_1} \cdot (C_A \cdot L_2)_{x_2} \cdot \text{mult}_{\text{ev}}(C_A) \cdot (C_A \cdot C_B)_{x_{n+1}=x_{n+2}} \cdot \text{mult}_{\text{ev}}(C_B) \end{aligned}$$

which completes the proof of case 2. We assumed that all edges e_1, \dots, e_4 are bounded. If one of the edges e_1 and e_2 (and/or one of the edges e_3 and e_4) is unbounded, the proof is essentially the same. If for example e_2 is unbounded, we do not have the l_2 -column, but there are also no marked points that can be reached from V via e_2 . \square

Of course there are completely analogous statements to lemma 7.17, remark 7.18, and proposition 7.21 if the $\mathcal{M}_{\text{trop}, 0,4}$ -coordinate of the curves in question is of type (2) instead of type (1) (that is, the two marked ends x_1 and x_3 come together at a 3-valent vertex — see example 4.9). However, there are no curves such that x_1 and x_2 are adjacent to the same vertex (as in part 1 of lemma 7.17) in this case, because x_1 and x_3 would have to map to $L_1 \cap p_3$, which is empty. Note that this is analogous to the classical case: there we had the two divisors $D_{1,2/3,4}$ and $D_{1,3/2,4}$ which (after intersecting them with the evaluations) correspond to stable maps where the underlying curve has a node. The two components of the underlying curves are mapped with degree d_A and d_B to \mathbb{P}^2 , leading to reducible image curves of degree d_A respectively d_B . In $D_{1,2/3,4}$ it is possible that $d_A = 0$ — the corresponding component with x_1 and x_2 is then mapped to the intersection point of the two lines which x_1 and x_2 are required to meet. The other component leads to a degree $d_B = d$ image curve. In $D_{1,3/2,4}$ neither d_A nor d_B can be 0, because x_1 and x_3 (respectively, x_2 and x_4) cannot be mapped to the same point, because $L_1 \cap p_3$ (respectively, $L_2 \cap p_4$) is empty (see section 3.2).

We are now ready to prove theorem 7.1. Note that the idea of this last proof is also the same as in the classical situation: all combinatorial factors we have to count arise for the same reasons — for the possibilities to arrange the points x_5, \dots, x_n on C_A and C_B , and for the possibilities that C_A intersects C_B or the two lines.

Proof of theorem 7.1:

We compute the degree of the map π of definition 7.12 at two different points. First consider a point $\mathcal{P}_1 = (a, b, p_3, \dots, p_n, C') \in \mathbb{R}^{2n-2} \times \mathcal{M}_{\text{trop}, 0,4}$ in π -general position with $\mathcal{M}_{\text{trop}, 0,4}$ -coordinate C' of type (1) (that is, x_1 and x_2 come together at a 3-valent vertex — see example 4.9) with a very large length. We have to count the curves in $\pi^{-1}(\mathcal{P}_1)$ with their respective π -multiplicity. Starting with the curves of type 1 in lemma 7.17 we see by proposition 7.21 that they count curves of degree d through $3d - 1$ points with their ev-multiplicity. Recall that lemma 4.68 shows in the case of rational curves that $\text{mult}_{\text{ev}} C = \text{mult} C$ (see remark 7.7). So these curves are counted with their ordinary multiplicity and give a contribution of $N_{\text{trop}}^{\text{irr}}(d, 0)$. For the curves of type 2 remark 7.18 tells us that we can as well count tuples $(C_A, C_B, x_1, \dots, x_n, x_{n+1}, x_{n+2})$, where

- (1) C_A and C_B are tropical curves of degrees d_A and d_B with $d_A + d_B = d$;
- (2) x_1, x_2 are marked points on C_A that map to L_1 and L_2 , respectively;
- (3) x_3, x_4 are marked points on C_B that map to p_3 and p_4 , respectively;

- (4) x_5, \dots, x_n are marked points that map to p_5, \dots, p_n and of which exactly $3d_A - 1$ lie on C_A and $3d_B - 1$ on C_B ;
- (5) $x_{n+1} \in C_A$ and $x_{n+2} \in C_B$ are points with the same image in \mathbb{R}^2 ;

where each such tuple has to be counted with the multiplicity computed in proposition 7.21.

Fix d_A and d_B with $d_A + d_B = d$. There are $\binom{3d-4}{3d_A-1}$ choices to split up the points x_5, \dots, x_n as in (4). Then we have $N_{\text{trop}}^{\text{irr}}(d_A, 0) \cdot N_{\text{trop}}^{\text{irr}}(d_B, 0)$ choices for C_A and C_B in (1). (As above, we have to count them with their multiplicity which is by 4.68 equal to their ev-multiplicity. The ev-multiplicity is the correct factor with which we have to count them due to proposition 7.21).

By Bézout’s theorem (see 7.20) there are d_A possibilities for x_1 in (2) — namely the intersection points of C_A with L_1 — if we count each of them with its local intersection multiplicity $(C_A \cdot L_1)_{x_1}$ as required by proposition 7.21. In the same way there are again d_A choices for x_2 and $d_A \cdot d_B$ choices for the gluing point $x_{n+1} = x_{n+2}$ as in (5). (Note that we can apply Bézout’s theorem without problems since we have seen in remark 7.16 that C_A intersects L_1, L_2 , and C_B in only finitely many points.)

Altogether the degree of π at \mathcal{P}_1 is

$$\deg_{\pi}(\mathcal{P}_1) = N_{\text{trop}}^{\text{irr}}(d, 0) + \sum_{\substack{d_A+d_B=d \\ d_A, d_B>0}} d_A^3 d_B \binom{3d-4}{3d_A-1} N_{\text{trop}}^{\text{irr}}(d_A, 0) N_{\text{trop}}^{\text{irr}}(d_B, 0).$$

Repeating the same arguments for a point \mathcal{P}_2 with $\mathcal{M}_{\text{trop}, 0,4}$ -coordinate of type (2) (that is, x_1 and x_3 come together at a 3-valent vertex as in example 4.9) we get

$$\deg_{\pi}(\mathcal{P}_2) = \sum_{\substack{d_A+d_B=d \\ d_A, d_B>0}} d_A^2 d_B^2 \binom{3d-4}{3d_A-2} N_{\text{trop}}^{\text{irr}}(d_A, 0) N_{\text{trop}}^{\text{irr}}(d_B, 0).$$

Equating these two expressions by proposition 7.14 now gives the desired result. □

8. THE TROPICAL CAPORASO-HARRIS ALGORITHM

In section 3.3 we described a way to compute the numbers $N_{\text{cplx}}(d, g)$ (which did not only work for rational curves as Kontsevich's formula) — the algorithm of Caporaso and Harris. In this recursion there are other numbers which we have to take into account: the numbers $N_{\text{cplx}}^{\alpha, \beta}(d, g)$ (defined in 3.35) of complex curves which satisfy tangency conditions (of higher order) to a given line in addition to passing through a set of given points. (Example 3.33 shows why we have to deal with these numbers in the recursion process.) In this chapter, we want to describe the tropical analogue of the Caporaso-Harris algorithm. That is, we want to prove that the numbers $N_{\text{trop}}(d, g)$ satisfy the same recursion formula. Before we can prove this tropically, we have to define tropical analogues for the numbers $N_{\text{cplx}}^{\alpha, \beta}(d, g)$, that is, we have to count tropical curves which satisfy tangency conditions (of higher order) to a given line. This will be done in section 8.1. The main theorem of this chapter — the theorem that the numbers $N_{\text{trop}}^{\alpha, \beta}(d, g)$ satisfy the Caporaso-Harris formula can of course only be formulated after these numbers have been defined (see theorem 8.6). It will be proved in section 8.2. An advantage of the tropical proof will be that we do not need to work with a moduli space. Instead, we can count the possibilities for the tropical curves satisfying our conditions purely combinatorially.

8.1 Remark

Note that analogously to remark 7.2, theorem 3.38 (which states that the numbers $N_{\text{cplx}}^{\alpha, \beta}(d, g)$ satisfy the Caporaso-Harris formula) together with theorem 8.6 (which states the same for the tropical numbers) can be used to give an alternative proof of Mikhalkin's Correspondence Theorem 6.1: knowing that both the numbers $N_{\text{cplx}}(d, g)$ and $N_{\text{trop}}(d, g)$ satisfy the same recursion formula, and knowing that there is one tropical as well as one complex line through two points, we can conclude that $N_{\text{cplx}}(d, g) = N_{\text{trop}}(d, g)$. We can even prove a generalized Correspondence Theorem, which states that also the numbers $N_{\text{cplx}}^{\alpha, \beta}(d, g)$ and $N_{\text{trop}}^{\alpha, \beta}(d, g)$ coincide. This proof of the Correspondence Theorem has of course the disadvantage that it does not give an intuitive idea why the two numbers should coincide, and also, it does not give a bijection from one set of curves to the other. However, at least the tropical side of this proof is much simpler as we will see in section 8.2. (But the classical proof of the Caporaso Harris algorithm is complicated.)

Beyond this implicit proof that $N_{\text{cplx}}^{\alpha, \beta}(d, g) = N_{\text{trop}}^{\alpha, \beta}(d, g)$ we want to give a direct proof of the generalized Correspondence Theorem analogous to Mikhalkin's proof that we described in chapter 6. The generalized Correspondence Theorem will be formulated and proved in a direct way in section 8.4.

In remark 3.51 we have seen that Caporaso and Harris also gave a recursion formula for irreducible curves. We can prove that the same recursion holds for the tropical numbers, too. This will be done in section 8.3.


In the whole chapter, we choose the point configurations in restricted general position. Therefore, we only have to deal with simple parametrized curves. Instead, we will work with unparametrized tropical curves. (We know that we can count these instead of the parametrized tropical curves by 5.34.)

Note that the methods used here do not necessarily require that we work with tropical curves of degree d . The results of this chapter can for example be generalized to tropical curves dual to a $d \times d'$ -rectangle - that is to tropical curves which come from complex curves on $\mathbb{P}^1 \times \mathbb{P}^1$. For the recursive formula for curves on $\mathbb{P}^1 \times \mathbb{P}^1$, see remark 9.14. In order to generalize theorem 8.6 to curves dual to other polygons, the combinatorial possibilities for those curves have to be studied in more detail.

The results of this chapter were achieved in joint work with Andreas Gathmann and published as preprint in [12].

8.1. TROPICAL CURVES THAT SATISFY HIGHER ORDER TANGENCY CONDITIONS TO A LINE

Before we define the numbers $N_{\text{trop}}^{\alpha,\beta}(d, g)$ we want to give a motivation of our definition. These numbers shall count the numbers of tropical curves which satisfy tangency conditions to a given line. The notion of tangency does not seem to have an easy analogue in the tropical world. Of course, we defined the tropical intersection multiplicity of two tropical curves in 7.19, and we could have the idea to define that a tropical curve is tangent to a line if this intersection multiplicity is bigger 1 at the intersection point. However, the following picture shows that then it would in general not be possible to draw a line tangent to a curve C in a given point $p \in C$. The picture shows two different tropical lines through p - none intersects with higher order.



If this point lies on an edge of direction $(-1, 0)$ as in the picture, then no tropical line intersects it with a higher intersection multiplicity. Also, if this point lies on an edge of direction $(2, 0)$, then every tropical line through p intersects with multiplicity 2 - so every line through p would be tangent. A solution for this problem might be to extend the definition of intersection multiplicity to intersection points which lie on a vertex of one of the curves.

Here, we want to work with another notion of tangency. The idea why we choose it is motivated by complex curves: Recall that our aim is to count complex curves with the aid of tropical geometry. As in chapter 6, we would like to count the images of complex curves under the map Log and the limit process instead of the complex curves themselves. So we have to find out what happens to a complex curve which satisfies tangency conditions to a given line under the map Log and the limit process. It is easier to find this out if we choose a special line, namely a coordinate line. As an example, let C be a complex conic which has a point of contact order 2 to the line $\{z = 0\}$. (As in chapter 6, notation 6.2, let (z, w) denote the coordinates of \mathbb{C}^2 and (x, y) the coordinates of \mathbb{R}^2 .) That is, C intersects the

line only in one point p . We have seen in the beginning of chapter 2 that each intersection point of C with $\{z = 0\}$ corresponds to a tentacle of the amoeba to the left — if we move on C towards the intersection point with $\{z = 0\}$, the coordinate of the image under Log tends to $-\infty$. (We have also seen an example for this in 6.11.) With the amoeba of C , also the spine of the amoeba has only one unbounded edge to the left. If we take the limit — that is, if we pass to J_t -holomorphic curves and let $t \rightarrow \infty$ — all spines of these J_t -holomorphic curves have only one unbounded end to the left. Hence also the tropical curve to which C is projected, the limit of these spines, has only one unbounded edge to the left. But as the tropical curve is of degree 2, this unbounded edge must have weight 2. So our idea is that tangency conditions (of higher order) to the special coordinate line $\{z = 0\}$ correspond to ends of a higher weight to the left.

This motivates the following definition.

Recall the notations from 3.34.

8.2 Definition

Let $d \geq 0$ and g be integers, and let α and β be sequences with $I\alpha + I\beta = d$. Let C be a simple (unparametrized) tropical curve of genus g and degree

$$\{(\alpha_i + \beta_i) \cdot (-i, 0), d \cdot (0, -1), d \cdot (1, 1) \mid i \in \mathbb{N}\}.$$

That is, C has $\alpha_i + \beta_i$ unbounded ends to the left of weight i for all i . Let α_i of these unbounded edges have a fixed position, that is, their y -coordinate is fixed. (We can think of this for example as follows: if we take the unique simple parametrization of C (see 5.29), then we can add a marked point adjacent to each of the corresponding ends and require this marked point to meet an image point with the prescribed y -coordinate and which lies far left of all other points.)

We define the (α, β) -multiplicity of C to be

$$\text{mult}_{\alpha, \beta}(C) := \frac{1}{I\alpha} \cdot \text{mult}(C)$$

where $\text{mult}(C)$ is the usual multiplicity as in definition 4.47.


Furthermore, define $N_{\text{trop}}^{\alpha, \beta}(d, g)$ to be the number of tropical curves of genus g and degree

$$\{(\alpha_i + \beta_i) \cdot (-i, 0), d \cdot (0, -1), d \cdot (1, 1) \mid i \in \mathbb{N}\}$$

with α_i fixed and β_i non-fixed unbounded ends to the left of weight i for all i that pass in addition through a set \mathcal{P} of $2d + g + |\beta| - 1$ points in restricted general position. The curves have to be counted with their respective (α, β) -multiplicities.


8.3 Example

The following picture shows a curve of degree $\{(-2, 0), (-1, 0), 3 \cdot (0, -1), 3 \cdot (1, 1)\}$. The end of weight 1 to the left is fixed — we have indicated this in the picture by a point. There are 6 more points through which the curve passes. As we have one fixed end of weight 1 and one nonfixed of weight 2, the sequences α and β are $\alpha = (1)$ and $\beta = (0, 1)$. The (α, β) -multiplicity of C is equal to the multiplicity of C .



8.4 Remark

Note that dual to the unbounded ends of such a curve is the Newton polygon Δ_d (see definition 3.69). For the curve in example 8.3, we have for example the following Newton polygon:



The ends of higher weight to the left correspond to steps of higher integer length in the boundary of the Newton polygon (see definition 5.15). Still, we have the same Newton polygon as for a curve of degree d . Also, the sum of the directions with primitive integral vector $(-1, 0)$ is $d \cdot (-1, 0)$, because we assumed that $I\alpha + I\beta = d$. Therefore, we will by abuse of notation speak of curves of degree d , even though there are ends of higher weight. This justifies also the notation $N_{\text{trop}}^{\alpha, \beta}(d, g)$ for the numbers of curves of degree

$$\{(\alpha_i + \beta_i) \cdot (-i, 0), d \cdot (0, -1), d \cdot (1, 1) \mid i \in \mathbb{N}\}$$

and genus g with the right ends and passing through the prescribed points.

8.5 Remark

Analogously to theorem 4.53 one can show that the numbers $N_{\text{trop}}^{\alpha, \beta}(d, g)$ (respectively, $N_{\text{trop}}^{\text{irr}, (\alpha, \beta)}(d, g)$) do not depend on the choice of fixed and nonfixed unbounded edges and points through which they shall pass. We have to evaluate only the y -coordinates of the unbounded edges which shall be fixed (together with the remaining marked points) then.

In the following section we will see that the definition of the numbers $N_{\text{trop}}^{\alpha, \beta}(d, g)$ is appropriate for our means: we will show that these numbers satisfy the Caporaso-Harris formula. As we have seen in remark 8.1, this tells us that they are equal to the corresponding complex numbers.

8.2. THE TROPICAL CAPORASO-HARRIS ALGORITHM

Having defined the numbers $N_{\text{trop}}^{\alpha, \beta}(d, g)$ we now want to prove the central result of this chapter. Note that the main idea of the proof is analogous to the classical case: there we specialized one point to lie on the line L to which also the tangency conditions are imposed (see section 3.3). Here, the complex line $\{z = 0\}$ that we choose is mapped by Log to the far left of the real plane. Hence, we will also move a point to the far left. A difference to the classical proof is that the curves we are counting do not really split into

two components. (Also, there is not a contracted bounded edge that we can cut as we have done it in remark 7.16 to interpret the cut curve as reducible. In fact, we decided to work with unparametrized tropical curves here which come from simple parametrized curves, so we do not have contracted bounded edges at all.) Having moved a point to the far left, the curves through the new point configuration can still be irreducible. However, we can consider two parts, one which corresponds to a line and which lies on the far left, and one on the right which will be a tropical curve of a lower degree.

8.6 Theorem


The numbers $N_{\text{trop}}^{\alpha,\beta}(d, g)$ defined in 8.2 satisfy the Caporaso-Harris formula (see definition 3.37).

Proof:

First, we choose a special point configuration through which the curves are required to pass. (As always, we require that the configuration is in restricted general position. We know that the numbers $N_{\text{trop}}^{\alpha,\beta}(d, g)$ do not depend on the chosen point configuration by 8.5.) Let $\varepsilon > 0$ be a small and $N > 0$ a large real number. We choose the fixed unbounded left ends and the set $\mathcal{P} = \{p_1, \dots, p_n\}$ so that

- the y -coordinates of all p_i and the fixed ends are in the open interval $(-\varepsilon, \varepsilon)$;
- the x -coordinates of p_2, \dots, p_n are in the open interval $(-\varepsilon, \varepsilon)$;
- the x -coordinate of p_1 is less than $-N$.

That is, we keep all points in a small horizontal strip and move p_1 very far to the left.



Let us consider an (unparametrized) tropical curve C satisfying the given conditions. We want to show that C must “look as in the picture above”, that is, that the curve “splits” into two parts joined by only horizontal lines. As we chose our point configuration in restricted general position, we can conclude that C is simple, and that both $\text{mult } C$ and the (α, β) -multiplicity of C are nonzero. By definition 4.47 it follows that the unique parametrization of C is rigid (see 5.29). Therefore, we know that there is no string. In particular, there is no path from one unbounded edge to another without meeting a marked point (see 4.46). Of course, also in the unparametrized tropical curve C there is no such path. (Note that a “crossing” of the images of two edges of the parametrized tropical curve is no connection.) We will use this in the following several times.

First we want to show that no vertex of C can have its y -coordinate below $-\varepsilon$. To see this, assume V is a vertex with lowest y -coordinate, and assume it is below $-\varepsilon$. By the balancing condition there must be an edge pointing downwards from V . As there is no vertex below V this must be an unbounded edge. As we prescribed the degree we know that this end must have direction $(0, -1)$ (especially, weight 1), and it must be the only edge pointing downwards. By the balancing condition it follows that at least one other edge starting at V must be horizontal: If not, both direction vectors of the other two edges adjacent to V would have a nonzero y -coordinate, and so the sum of these two directions could not be equal to $(0, 1)$ which is necessary by the balancing condition. Along this horizontal edge we can go (again due to the balancing condition) from V to another unbounded edge in the region $\{y \leq -\varepsilon\}$. As there are no marked points in this region this means that there are two unbounded edges that we can reach from V without meeting a marked point (the one pointing downwards from V and the horizontal one). This is a contradiction to the above as C must be rigid.

Analogously, we can see that no vertex of C can have its y -coordinate above ε .

Next, consider the rectangle

$$R := \{(x, y); -N \leq x \leq -\varepsilon, -\varepsilon \leq y \leq \varepsilon\}.$$

We want to study whether there can be vertices of C within R . Let C_0 be a connected component of $C \cap R$. Note that any edge of C_0 leaving R at the top or bottom edge must be unbounded as we have just seen that there are no vertices of C above or below R . If there are edges of C_0 leaving R at the top *and* at the bottom then we could again go from one unbounded edge of C to another without passing a marked point, which is not possible by the above. So we can assume without loss of generality that C_0 does not meet the top edge of R .

Again due to the balancing condition the edges of $C_0 \cap R$ that are not horizontal must then project to the x -axis to a union of two (maybe empty) intervals $[-N, x_1] \cup [x_2, -\varepsilon]$. (Otherwise, there would also be an edge which meets the top edge of R .) But the number of edges of C as well as the minimum slope of an edge (and hence the maximum distance an edge can have within R) are bounded by a constant that depends only on the degree d of the curves. So we can find $a, b \in \mathbb{R}$, $-N \leq a < b \leq -\varepsilon$ (that depend only on d) such that the interval $[a, b]$ is disjoint from $[-N, x_1] \cup [x_2, -\varepsilon]$, or in other words such that for any curve C which satisfies the conditions, there are no non-horizontal edges in $[a, b] \times [-\varepsilon, \varepsilon]$. Also, there are then no vertices of C in $[a, b] \times \mathbb{R}$. Hence we can see that the curve C must look as in the picture above: we can “cut” it at any line $x = x_0$ with $a < x_0 < b$ and obtain curves on both sides of this line that are joined only by horizontal lines.

There are now two cases to distinguish, corresponding to the two types of summands which appear in the Caporas-Harris formula:

- (1) p_1 lies on a horizontal non-fixed end of C . Then the region where $x \leq -\varepsilon$ consists of only horizontal lines. (If not, there would be at least two unbounded edges of direction $(0, -1)$ and $(1, 1)$ in this region, and we would have a path from one to the other without meeting a marked point.) We can hence consider C as having

one more fixed end at p_1 and passing through $\mathcal{P} \setminus \{p_1\}$. We just have to multiply with the weight of this end, as the multiplicity of curves with fixed ends is defined as $\frac{1}{I^\alpha} \cdot \text{mult}(C)$ by 8.2. Therefore the contribution of these curves to $N_{\text{trop}}^{\alpha, \beta}(d, g)$ is

$$\sum_{k: \beta_k > 0} k \cdot N_{\text{trop}}^{\alpha+e_k, \beta-e_k}(d, g).$$

Note that this is the first type of summands that occur in the Caporaso-Harris formula.

- (2) p_1 does not lie on a horizontal end of C (as in the picture above). Then there must also be unbounded edges of direction $(0, -1)$ and $(1, 1)$ in this region. Due to the balancing condition we must have as many of direction $(0, -1)$ as of $(1, 1)$. But as there is only one marked point — namely p_1 — to separate the unbounded edges in this region, there cannot be more than two. So the left part has exactly one end in direction $(0, -1)$ and $(1, 1)$ each, together with some more horizontal ends. Hence the curve on the right must have degree $d - 1$. Let us denote this curve by C' .

How many possibilities are there for C' ? Assume that $\alpha' \leq \alpha$ of the fixed horizontal ends only intersect the part $C \setminus C'$ and are not adjacent to a 3-valent vertex of $C \setminus C'$. Then C' has α' fixed horizontal ends. There are $\binom{\alpha}{\alpha'}$ possibilities for this case. C' has $d - 1 - I\alpha'$ non-fixed horizontal ends. Let β' be a sequence which fulfills $I\beta' = d - 1 - I\alpha'$, hence a possible choice of weights for the non-fixed ends of C' . Assume that $\beta'' \leq \beta'$ of these ends are adjacent to a 3-valent vertex of $C \setminus C'$ whereas $\beta' - \beta''$ ends intersect $C \setminus C'$, that is, just cross. The connected component of $C \setminus C'$ which contains p_1 has to contain the two ends of direction $(0, -1)$ and $(1, 1)$ due to the balancing condition. (Recall that connected component means: the image of a connected component of a parametrization of smallest genus.) Also, it can contain some ends of direction $(-1, 0)$ — but this have to be fixed ends then, as p_1 cannot separate more than two (nonfixed) ends. So all the β nonfixed ends of direction $(-1, 0)$ have to intersect $C \setminus C'$ - therefore they have to be ends of C' . That is, $\beta' - \beta'' = \beta$. There are $\binom{\beta'}{\beta}$ possibilities for this. Furthermore, we have

$$\begin{aligned} \text{mult}_{\alpha, \beta}(C) &= \frac{1}{I^\alpha} \text{mult}(C) = \frac{1}{I^\alpha} \cdot I^{\alpha-\alpha'} \cdot I^{\beta'-\beta} \cdot \text{mult}(C') \\ &= \frac{1}{I^{\alpha'}} \cdot I^{\beta'-\beta} \cdot \text{mult}(C') = I^{\beta'-\beta} \cdot \text{mult}_{\alpha', \beta'}(C') \end{aligned}$$

where the factors $I^{\alpha-\alpha'}$ and $I^{\beta'-\beta}$ arise due to the 3-valent vertices which are not part of C' . The genus g' of C' is equal to $g' = g - (|\beta' - \beta| - 1)$ since we cut $|\beta' - \beta| - 1$ loops as in the picture above (at least if C and C' are both irreducible. If not, the number of loops is equal to $|\beta' - \beta| - 1$ minus the number of components of C' , but the genus of C' is the number of loops of C' minus the number of components plus 1 too by remark 4.4, and therefore the equation is still fulfilled). Furthermore, $g - g' \leq d - 2$ as at most $d - 2$ circles may be cut.

The sum of these two contributions gives the required Caporaso-Harris formula. \square

Note that similarly to the tropical proof of Kontsevich’s formula (see chapter 7), the combinatorial factors that arise in this formula here arise due to similar reasons as in the classical proof. For example, we get the factor $\binom{\alpha}{\alpha'}$, because we have as many possibilities to have fixed unbounded edges (that is, fixed tangency conditions) on the curve C' . In the classical case, we get the same factor, because there are as many possibilities to arrange the marked points which shall satisfy the tangency conditions on the components of the curve.

Also, in the classical case, we get the factor $I^{\beta'-\beta}$, because the corresponding component in the moduli space occurs with the higher multiplicity $I^{\beta'-\beta}$ (see 3.48). Here the “corresponding component” — that is, the set of all curves which have β' ends of which β just intersect $C \setminus C'$ — has to be counted with the factor $I^{\beta'-\beta}$, because this is the multiplicity of the left part $C \setminus C'$, to which the $\beta' - \beta$ ends are adjacent. So, here the argument is purely combinatorial, contrary to the classical proof. That is, using the generalized Correspondence Theorem 8.9, we can give an alternative and much easier proof of theorem 3.38 stating that the complex numbers $N_{\text{trop}}^{\alpha,\beta}(d, g)$ satisfy the Caporaso Harris formula.

8.3. THE TROPICAL CAPORASO-HARRIS ALGORITHM FOR IRREDUCIBLE CURVES

As mentioned in remark 3.51, Caporaso and Harris also gave an algorithm to compute the numbers $N_{\text{cplx}}^{\text{irr},(\alpha,\beta)}(d, g)$ of irreducible complex curves satisfying tangency conditions (of higher order) to a line in addition to passing through some given points. So far, we worked with unparametrized tropical curves, and we know that they can be reducible if they allow a parametrization where the underlying graph is not connected. Here, we want to prove that also the irreducible Caporaso-Harris formula is satisfied by the analogous tropical numbers. The proof is analogous to the proof of theorem 8.6 — we will also choose the special point configuration and divide the curve into two parts. However, it is harder to control that we actually count only irreducible curves when we cut the curve into two parts. Note that an unparametrized curve C as in the proof of theorem 8.6 is irreducible if and only if every connected component of C' is linked to $C \setminus C'$. That is, for every connected component there has to be an unbounded edge of C' which does not only pass through $C \setminus C'$, but is adjacent to a 3-valent vertex. Also, as we want to construct a recursive formula for the numbers of irreducible curves, we cannot just count the possibilities for the curve C' , because it might be reducible. Instead, we have to count the connected components of C' separately. (Note that by connected components of C' we denote the connected components of a graph of least possible genus which parametrizes C' .)

8.7 Definition

Let $N_{\text{trop}}^{\text{irr},(\alpha,\beta)}(d, g)$ be the *number of irreducible (simple unparametrized) tropical curves of degree of genus g and degree*

$$\{(\alpha_i + \beta_i) \cdot (-i, 0), d \cdot (0, -1), d \cdot (1, 1) \mid i \in \mathbb{N}\}$$

with α_i fixed and β_i non-fixed horizontal left ends of weight i for all i that pass in addition through a set of $|\beta| + 2d + g - 1$ points in general position. Again, the curves have to be counted with their (α, β) -multiplicity as in definition 8.2.

Analogously to remark 8.5 we can conclude that these numbers do not depend on the special choice of fixed ends and points.

Recall that irreducible means here: the graph of a parametrization of least possible genus is connected (see definition 5.6).

8.8 Theorem

The numbers $N_{\text{trop}}^{\text{irr},(\alpha,\beta)}(d, g)$ satisfy the recursion relations from remark 3.51 known for the numbers $N_{\text{cplx}}^{\text{irr},(\alpha,\beta)}(d, g)$.

Proof:

The proof is analogous to that of theorem 8.6. Fix a set $\mathcal{P} = \{p_1, \dots, p_n\}$ with p_1 very far left as in the proof of theorem 8.6. Let C be an irreducible tropical curve with the right properties through the points. The first term in the recursion formula (that corresponds to curves with only horizontal lines in the area where $x \leq -\varepsilon$) follows in the same way as in theorem 8.6. So assume that p_1 does not lie on a horizontal end of C . As before we get a curve C' of degree $d-1$ to the right of the cut. The curve C' does not need to be irreducible however. That is, the graph of a parametrization of least possible genus is not connected and decomposes into k connected components $\Gamma_1, \dots, \Gamma_k$. As the balancing condition must be fulfilled for each connected component, we can conclude that $C_i = h|_{\Gamma_i}(\Gamma_i)$ is a curve of degree d_i . (As before, by abuse of notation we speak of degree d_i here even if there are ends of a higher weight, see 8.4.) Then $d_1 + \dots + d_k = d - 1$. How many possibilities are there for C_1, \dots, C_k ? Assume that α^j of the fixed horizontal ends only intersect the part $C \setminus C'$ and are not adjacent to a 3-valent vertex of $C \setminus C'$, but are adjacent to a 3-valent vertex of C_j . Then the irreducible tropical curve C_j has α^j fixed horizontal ends. Furthermore, $\alpha^1 + \dots + \alpha^k \leq \alpha$, and there are $\binom{\alpha}{\alpha^1, \dots, \alpha^k}$ possibilities for this. Let C_j have β^j more horizontal ends satisfying $|\beta^j| + |\alpha^j| = d_j$. Then C_j is fixed by $2d_j + g_j + |\beta^j| - 1$ points, where g_j denotes the genus of C_j . Therefore, there are

$$\binom{2d + g + |\beta| - 2}{2d_1 + g_1 + |\beta^1| - 1, \dots, 2d_k + g_k + |\beta^k| - 1}$$

possibilities to distribute the points p_2, \dots, p_n on the components C_1, \dots, C_k . Assume that $\beta^{j'}$ of the β^j ends are adjacent to a 3-valent vertex of $C \setminus C'$ whereas $\beta^j - \beta^{j'}$ ends just intersect $C \setminus C'$. There are $\binom{\beta^j}{\beta^j - \beta^{j'}}$ possibilities for this, and the product of the multiplicities of the 3-valent vertices of $C \setminus C'$ is $I^{\beta^{j'}}$. As C is irreducible we must have $|\beta^{j'}| > 0$: That is, each component C_j has to be linked to $C \setminus C'$ at a 3-valent vertex. It is not possible that each end of C_j just intersects $C \setminus C'$, because then C_j would be a component of C . But C is assumed to be irreducible. For each C_j we see that $|\beta^{j'}| - 1$ circles are cut (as in the picture in the proof of theorem 8.6). Hence,

$g = \sum g_j + \sum(|\beta^{j'}| - 1) = \sum g_j + \sum |\beta^{j'}| - k$. Furthermore,

$$\begin{aligned} 2d + g + |\beta| - 1 - 1 &= \sum 2d_j + g_j + |\beta^j| - 1 \\ \Rightarrow 2(d-1) + \sum g_j + \sum |\beta^{j'}| - k + |\beta| &= 2(d-1) + \sum g_j - k + \sum |\beta^j| \\ \Rightarrow |\beta| + \sum |\beta^{j'}| &= \sum |\beta^j| \\ \Rightarrow \beta + \sum \beta^{j'} &= \sum \beta^j. \end{aligned}$$

This proves the recursion formula except for the factor $\frac{1}{\sigma}$. This factor is needed because up to now we count different curves if two components C_i and C_j of C' are identical, depending on whether C_i is the i -th component or C_j is the i -th component. Therefore we have to divide by σ . \square

8.4. THE CORRESPONDENCE OF COMPLEX CURVES TANGENT (OF HIGHER ORDER) TO A LINE AND TROPICAL CURVES WITH ENDS OF HIGHER WEIGHT

Of course, the tropical proofs for both Caporaso-Harris algorithms (the one for not necessarily irreducible curves as well as the one for irreducible curves) together with Caporaso's and Harris' original proofs of the algorithms for the complex numbers tell us that we have $N_{\text{cplx}}^{\alpha,\beta}(d, g) = N_{\text{trop}}^{\alpha,\beta}(d, g)$ and $N_{\text{cplx}}^{\text{irr},(\alpha,\beta)}(d, g) = N_{\text{trop}}^{\text{irr},(\alpha,\beta)}(d, g)$ by using the recursion and the initial value of one complex as well as one tropical line through two given points (see remark 8.1). However, here we would like to give a more direct proof of the correspondence $N_{\text{cplx}}^{\alpha,\beta}(d, g) = N_{\text{trop}}^{\alpha,\beta}(d, g)$ and $N_{\text{cplx}}^{\text{irr},(\alpha,\beta)}(d, g) = N_{\text{trop}}^{\text{irr},(\alpha,\beta)}(d, g)$ that shows us which complex curves project to which tropical curves.

8.9 Theorem (“Generalized Correspondence Theorem”)

For the numbers $N_{\text{cplx}}^{\alpha,\beta}(\Delta, g)$ defined in 3.35, $N_{\text{cplx}}^{\text{irr},(\alpha,\beta)}(\Delta, g)$ defined in 3.51, $N_{\text{trop}}^{\alpha,\beta}(\Delta, g)$ defined in 8.2 and $N_{\text{trop}}^{\text{irr},(\alpha,\beta)}(\Delta, g)$ defined in 8.7 we have:

$$\begin{aligned} N_{\text{trop}}^{\alpha,\beta}(d, g) &= N_{\text{cplx}}^{\alpha,\beta}(d, g) \text{ and} \\ N_{\text{trop}}^{\text{irr},(\alpha,\beta)}(d, g) &= N_{\text{cplx}}^{\text{irr},(\alpha,\beta)}(d, g) \end{aligned}$$

for all d, g, α, β .

The proof we are going to present here is related to the proof of Mikhalkin's Correspondence Theorem 6.1. Let \mathcal{Q} be a set of $n = |\beta| + 2d + g - 1$ points in general position in $(\mathbb{C}^*)^2$ such that $\mathcal{P} = \text{Log}(\mathcal{Q})$ is in restricted tropical general position in \mathbb{R}^2 (defined in 5.33). We know that there are finitely many parametrized tropical curves of genus g and degree

$$\{(\alpha_i + \beta_i) \cdot (-i, 0), d \cdot (0, -1), d \cdot (1, 1) \mid i \in \mathbb{N}\}$$

through \mathcal{P} , where α_i of the unbounded edges to the left are fixed, and each has a nonzero multiplicity. Call these tropical curves C_1, \dots, C_r . Note that by lemma 5.34 it makes no difference whether we consider the C_j as parametrized tropical curves or as unparametrized tropical curves.

First, we generalize lemma 6.6:

8.10 Lemma

If t is large enough then for all J_t -holomorphic curves V of genus g and degree d passing through \mathcal{Q} and with contact order i with the line $\{z = 0\}$ in α_i fixed and β_i arbitrary points the amoeba $\text{Log}(V)$ is contained in a small neighborhood of C_j for some j .

The proof of this lemma follows from 8.12 and the following considerations. In order to prove this lemma we only need to see that the amoebas of J_t -holomorphic curves with the right properties are close to tropical curves with α_i fixed and β_i non-fixed horizontal ends of weight i . With the help of the proof of 6.6 it then follows that the tropical curves are of the right genus and degree and pass through \mathcal{P} (and must therefore be one of the C_j). If $(V_k)_{k \in \mathbb{N}}$ is a sequence of curves with the right properties and such that V_k is J_{t_k} -holomorphic ($t_k \rightarrow \infty$) then C_j is actually the limit of the spines of a subsequence of V_k (see 6.15). Therefore we only need to see that the spines of the amoebas of J_t -holomorphic curves with the right properties have α_i fixed and β_i non-fixed horizontal ends of weight i . This will be shown in 8.12.

In 6.9 we introduced the spine of an amoeba \mathcal{A} as a tropical curve S which lies inside the amoeba. In particular, every connected component of $\mathbb{R}^2 \setminus \mathcal{A}$ corresponds to a connected component of $\mathbb{R}^2 \setminus S$. That is, if we draw the dual Newton subdivision of the spine, there is an integer point in Δ_d (see definition 3.69) for each connected components of $\mathbb{R}^2 \setminus \mathcal{A}$. This integer point in Δ_d is called the order of the connected component of the complement of \mathcal{A} (for a precise definition, see [8]). Let V be a J_t -holomorphic curve. Then V is given by a polynomial $f(z, w)$. The contact order conditions for V at $\{z = 0\}$ tell us that there exists a polynomial g such that

$$f(z, w) = z \cdot g(z, w) + \prod_{j=1}^r (w - p_j)^{m_j}$$

where the $(0, p_i)$ are the points of higher contact order with the line $\{z = 0\}$, sorted so that $|p_1| < \dots < |p_r|$, and the m_j are the contact orders. The amoeba \mathcal{A}_f is given as $\text{Log}(\{f(z, w) = 0\})$. Let K be a connected component of the complement $\mathbb{R}^2 \setminus \mathcal{A}_f$. Then the order $\nu(K)$ is a lattice point in Δ_d . We are going to use the following lemma (2.2) of [8]:

8.11 Lemma

Let u be a point in a connected component K of the complement of an amoeba \mathcal{A}_f , $K \subset \mathbb{R}^2 \setminus \mathcal{A}_f$, choose c such that $\text{Log}(c) = u$, and denote by $\nu(K)$ the order of the complement containing u . For all integer vectors $v \in \mathbb{Z}^2$ with positive entries v_1, v_2 the canonical scalar product $\langle v, \nu(K) \rangle$ is equal to the number of zeros of the polynomial $w \mapsto f(c_1 w^{v_1}, c_2 w^{v_2})$ in the disc $|w| < 1$.

8.12 Lemma

Let V be a J_t -holomorphic curve with α_i fixed and β_i non-fixed points of contact order i with $\{z = 0\}$. Then the spine of the amoeba of V has α_i fixed and β_i non-fixed horizontal ends of weight i .

Proof:

As V intersects $\{z = 0\}$ in $|\beta| + |\alpha|$ points the amoeba has $|\beta| + |\alpha|$ tentacles in the

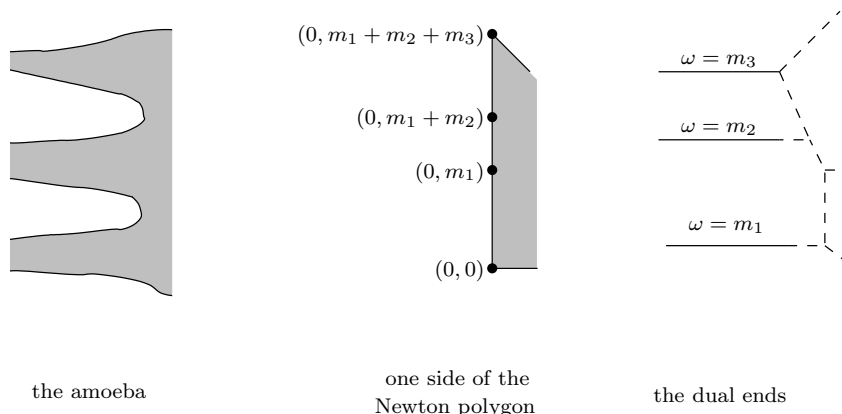
direction of $\text{Log}(\{z = 0\}) = \{x \rightarrow -\infty\}$, and the tentacles are close to $y = \log(|p_j|)$. Therefore there are $|\beta| + |\alpha| + 1$ connected components of the complement on the left part of the amoeba. We compute the order of each of these components with the help of lemma 8.11. Let K be the component between the tentacles close to $\log(|p_k|)$ and $\log(|p_{k+1}|)$. Let $c = (c_1, c_2) \in (\mathbb{C}^*)^2$ be a point such that $\text{Log}(c) \in K$. We can choose c_1 close to zero so that $\text{Log}(c)$ is a point on the far left. Furthermore, $|p_k| < |c_2| < |p_{k+1}|$. Let $v = (1, 0)$ and $v' = (0, 1)$. The number of zeros of $f(c_1 w^{v_1}, c_2 w^{v_2})$ in $|w| < 1$ is

$$\begin{aligned} \#\{\text{zeros of } f(c_1 w^{v_1}, c_2 w^{v_2})\} &= \#\{\text{zeros of } f(c_1 w, c_2)\} \\ &= \#\{\text{zeros of } c_1 w \cdot g(c_1 w, c_2) + \prod (c_2 - p_j)^{m_j}\} = \#\{\text{zeros of } \prod (c_2 - p_j)^{m_j}\} \end{aligned}$$

as c_1 is close to zero. But c_2 is not equal to any of the p_j , so this polynomial does not have any zeros, and therefore the first component of the order is zero. The number of zeros of $f(c_1 w^{v'_1}, c_2 w^{v'_2})$ is

$$\begin{aligned} \#\{\text{zeros of } f(c_1 w^{v'_1}, c_2 w^{v'_2})\} &= \#\{\text{zeros of } f(c_1, c_2 w)\} \\ &= \#\{\text{zeros of } c_1 \cdot g(c_1, c_2 w) + \prod (c_2 w - p_j)^{m_j}\} = \#\{\text{zeros of } \prod (c_2 w - p_j)^{m_j}\}. \end{aligned}$$

This polynomial has m_j zeros at $w = \frac{p_j}{c_2}$, and we have $|w| < 1$ if $j = 1, \dots, k$ and $|w| > 1$ if $j = k + 1, \dots, r$. So the polynomial has $m_1 + \dots + m_k$ zeros in $|w| < 1$. Hence the second component of the order is $m_1 + \dots + m_k$ by 8.11.



The picture above then shows that the dual tropical curve (that is the spine) has the right number of horizontal ends of weight i . Furthermore, if p_j was a fixed point then also the end at $y = \log(|p_j|)$ is fixed. □

Now, we generalize lemma 6.7. It shows that for each C_j there are $\frac{\text{mult}(C_j)}{\mu_{\text{edge}}(C_j, \mathcal{P})}$ complex tropical curves (see 6.18 and 6.22) projecting to C_j , and that for each of these complex tropical curves there are $\mu_{\text{edge}}(C_j, \mathcal{P})$ (see 6.24 and 6.25) J_t -holomorphic curves nearby. In order to generalize these ideas to our case we first need to generalize the definition of edge multiplicity.

8.13 Definition

Let $C = C_j$ be one of the tropical curves through \mathcal{P} . Think of $C = (\Gamma, h, x_1, \dots, x_n)$ as a parametrized tropical curve. Each edge of Γ has a certain weight which is equal to the factor with which we need to multiply the primitive integer vector u to get the direction

v (see remark 4.13). We define the *edge multiplicity* of C and \mathcal{P} , $\mu_{\text{edge}}(C, \mathcal{P})$, to be the product of the weights of all bounded edges.

If we think of C as an unparametrized tropical curve, then the edge multiplicity $\mu_{\text{edge}}(C, \mathcal{P})$ is equal to the product of the weights of all edges that are bounded and disjoint from \mathcal{P} or unbounded and not disjoint from \mathcal{P} , times the product of the squares of the weights of all edges that are bounded and not disjoint from \mathcal{P} .


The following proposition generalizes 6.22.

8.14 Proposition

Let $C = C_j$ be one of the tropical curves through \mathcal{P} . Then there are $\frac{\text{mult}(C)}{\mu_{\text{edge}}(C, \mathcal{P})}$ different complex tropical curves V (see 6.16) of the same degree and genus as C through \mathcal{Q} and such that $\text{Log}(V) = C$.

Proof:

The idea of the proof is the same as for 6.22. Let $(\Gamma, h, x_1, \dots, x_n)$ be a parametrization of C . We know that $\text{mult } C \neq 0$ as we chose \mathcal{P} in restricted general position. Therefore C is rigid and each connected component of $\Gamma \setminus \bigcup_i \overline{x_i}$ is rational and contains exactly one unbounded edge by 4.50. Let K' be one of these components. Let $p_i, p_j \in \mathcal{P}$ be the endpoints of two edges adjacent to the same 3-valent vertex V , and let $\Delta' \subset \Delta_d$ be the triangle dual to V .



Then 6.21 tells us that there are $\frac{2 \text{Area}(\Delta')}{\omega_i \omega_j}$ complex tropical curves such that their image under Log is dual to Δ' , where ω_i and ω_j denote the weights of the edges through p_i and p_j respectively. The result now follows by induction. □

Now analogously to Mikhalkin’s proof we have to see that for each of these complex tropical curves there are $\mu_{\text{edge}}(C, \mathcal{P})$ J_t -holomorphic curves with the required properties nearby. Recall that this is the most difficult part of the proof of the Correspondence Theorem. It is proved first for the case that there are no edges of a higher weight (see 6.24), and then separately for the case that there are edges of higher weight (see 6.25). Recall that by remark 6.26 it is essential for the proof of 6.25 that we allow only bounded edges of higher weight. Because we have also unbounded edges of higher weight here, it is not straightforward to generalize this idea for our case. Instead, we use the Correspondence Theorem for curves without ends of higher weight. To be able to use it we need to modify the Newton polygon such that there are no edges of integer length bigger than 1 in the boundary. The polygon Δ_d has $\beta_i + \alpha_i$ edges of length i at $\{x = 0\}$.

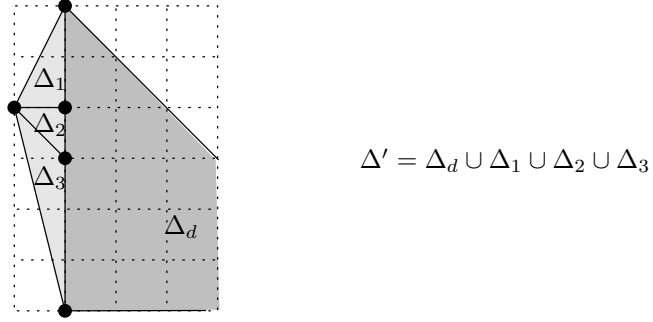
8.15 Lemma

Let C be one of the tropical curves C_j through \mathcal{P} . If t is large enough then for each of

the complex tropical curves V projecting to C there are exactly $\mu_{\text{edge}}(C, \mathcal{P})$ J_t -holomorphic curves of genus g and degree d through \mathcal{Q} and with the right contact orders to the line $\{z = 0\}$ in a small neighborhood of V .

Proof:

We enlarge Δ_d to a polygon Δ' by adding another point such that Δ' has no edges of higher integer length.



For this new polygon Δ' and for all tropical curves C' dual to Δ' we know that

$$\mu_{\text{edge}}(C', \mathcal{P}') = \mu_{\text{edge}}(C, \mathcal{P}) \cdot \prod (I^{\beta_i} \cdot I^{\alpha_i})$$

because we transferred the unbounded edges of weight i to bounded edges (\mathcal{P}' is \mathcal{P} plus another point condition on the new edges, such that we have finitely many tropical curves of Δ' , plus a point condition for each of the fixed ends). The multiplicity of C' is

$$\text{mult}(C') = \prod (I^{\beta_i} \cdot I^{\alpha_i}) \cdot \text{mult}(C)$$

as we add triangles of these areas. Therefore there are

$$\frac{\text{mult}(C')}{\mu_{\text{edge}}(C', \mathcal{P}')} = \frac{\text{mult}(C)}{\mu_{\text{edge}}(C, \mathcal{P})}$$

complex tropical curves projecting to C' . By 6.24 and 6.25 each of these complex tropical curves give rise to

$$\mu_{\text{edge}}(C', \mathcal{P}') = \mu_{\text{edge}}(C, \mathcal{P}) \cdot \prod (I^{\beta_i} \cdot I^{\alpha_i})$$

J_t -holomorphic curves of genus g and passing through \mathcal{Q}' in a neighborhood. These J_t -holomorphic curves arise as the gluing of little pieces corresponding to subpolygons of Δ' . If we set the gluing of the pieces corresponding to the new triangles of Δ' aside (that is, corresponding to $\Delta_1, \dots, \Delta_m$) we get J_t -holomorphic curves near the complex tropical curves projecting to C . Furthermore, these J_t -holomorphic curves have the right contact orders to the toric divisor corresponding to the left side $\{x = 0\}$ of Δ_d , as otherwise the genus would not be correct (see remark 6.27). So for each complex tropical curve V projecting to C we have $\mu_{\text{edge}}(C', \mathcal{P}')$ of these J_t -holomorphic curves with the right properties. However we count some of them too often: when gluing the pieces corresponding to the new triangles $\Delta_1, \dots, \Delta_m$ there are different possibilities to glue. Assume Δ_i has a side of integer length l on the left. Then we have l different ways of gluing the piece corresponding to Δ_i to the J_t -holomorphic curve near a complex tropical curve projecting

to C (see 6.27). Therefore, in order to get the number of different J_t -holomorphic curves near V we have to divide by $\prod(I^{\beta_i} \cdot I^{\alpha_i})$. So we get

$$\frac{\mu_{\text{edge}}(C', \mathcal{P}')}{\prod(I^{\beta_i} \cdot I^{\alpha_i})} = \mu_{\text{edge}}(C, \mathcal{P})$$

different J_t -holomorphic curves with the right properties in a neighborhood of V . \square

Analogously to the proof of theorem 6.1, we can sum up the arguments in order to prove 8.9.

9. THE CAPORASO-HARRIS ALGORITHM IN THE LATTICE PATH SETTING


In chapter 5 we have seen that there is a third way to determine the numbers $N_{\text{cplx}}(d, g)$ — we can also count lattice paths instead. Analogously to the previous chapter, our aim here is to reprove the algorithm of Caporaso and Harris in the lattice path setting, that is, for the numbers $N_{\text{path}}(d, g)$. Before we can do that, we need of course again a definition of $N_{\text{path}}^{\alpha, \beta}(d, g)$. We will deal with these numbers of *generalized lattice paths* in section 9.1. In fact, the definition is motivated by considering the duals of the tropical curves with ends of higher weights as we defined them in section 8.1. The theorem that the numbers $N_{\text{path}}^{\alpha, \beta}(d, g)$ satisfy the Caporaso-Harris formula (theorem 9.13) is again proved with completely different methods — it uses the definition of multiplicity of a lattice path and counts purely combinatorially the possibilities for our generalized lattice paths. This theorem is formulated and proved in section 9.2. The main idea is to count the possible Newton subdivisions for a path γ differently from how we have defined them in 5.40. In fact, we show that the number of possible Newton subdivisions is equal to the number of *column-wise Newton subdivisions*, a notion which will be made precise in section 9.2. These column-wise Newton subdivisions can now be counted column by column — therefore we get a recursive formula by dropping the first column and counting the possibilities in the triangle Δ_{d-1} corresponding to curves of lower degree.

Knowing that both the numbers $N_{\text{path}}^{\alpha, \beta}(d, g)$ and the numbers $N_{\text{trop}}^{\alpha, \beta}(d, g)$ satisfy the Caporaso-Harris formula, we can of course conclude by induction that $N_{\text{path}}^{\alpha, \beta}(d, g) = N_{\text{trop}}^{\alpha, \beta}(d, g)$ (and also $N_{\text{path}}^{\alpha, \beta}(d, g) = N_{\text{cplx}}^{\alpha, \beta}(d, g)$). However, we want to give a direct proof here, too, using the methods of chapter 5, where we cited Mikhalkin's proof of theorem 5.44 that $N_{\text{path}}(d, g) = N_{\text{trop}}(d, g)$. The theorem that $N_{\text{path}}^{\alpha, \beta}(d, g) = N_{\text{trop}}^{\alpha, \beta}(d, g)$ (see 9.15) will be proved in section 9.3. As there is no analogue to irreducible curves in the lattice path setting (see remark 5.45), we do not study the Caporaso-Harris algorithm for irreducible curves here.

The results of this chapter were achieved in joint work with Andreas Gathmann and published as preprint in [12].

9.1. GENERALIZED LATTICE PATHS

Our aim is now to slightly generalize the definitions of section 5.3 in order to allow more lattice paths and arrive at lattice path analogues of the numbers $N_{\text{cplx}}^{\alpha, \beta}(d, g)$. The idea what the general lattice path should look like arises after reconsidering the tropical curves that satisfy tangency conditions to a line as defined in 8.1. These were tropical curves with ends of higher weights to the left. Dual to such a tropical curve is a Newton subdivision of the triangle Δ_d (see definition 3.69), that does not contain all points of the left boundary. Instead, there have to be steps of bigger integer length:



Following the ideas of chapter 5, the tropical curves we count are dual to the possible Newton subdivisions of a path γ . The generalized paths that we need to allow here must therefore have the corresponding steps of bigger integer lengths at the left side of the boundary of Δ_d in their possible Newton subdivisions. This motivates the definition below.


9.1 Notation

For the whole chapter, choose the function $\lambda(x, y) = x - \varepsilon y$ where ε is a small number.

Choose two sequences α and β with $I\alpha + I\beta = d$. Let $\gamma : [0, n] \rightarrow \Delta_d$ be a λ -increasing path with $\gamma(0) = (0, d - I\alpha) = (0, I\beta)$ and $\gamma(n) = q = (d, 0)$. We are going to define a multiplicity for such a path γ . Again this multiplicity will be the product of a “positive” and a “negative” multiplicity that we define separately.

9.2 Definition

Let $\delta_\beta : [0, |\beta| + d] \rightarrow \Delta_d$ be a path such that the image $\delta_\beta([0, |\beta| + d])$ is equal to the piece of boundary of the triangle Δ_d (see definition 3.69) between $(0, I\beta)$ and $q = (d, 0)$, and such that there are β_i steps (i.e. images of a size one interval $[j, j + 1]$) of integer length i at the side $\{x = 0\}$ (and hence at $\{y = 0\}$ only steps of integer length 1). We define the negative multiplicity $\mu_{\beta,-}(\delta_\beta)$ of all such paths to be 1. For example, the following picture shows all paths δ_β for $\beta = (2, 1)$ and $d = 5$:



Using these starting paths the *negative multiplicity* $\mu_{\beta,-}(\gamma)$ of an arbitrary path as above is now defined recursively by the same procedure as in definition 5.39 (2).

9.3 Definition

To compute the *positive multiplicity* $\mu_{\alpha,+}(\gamma)$ we extend γ to a path $\gamma_\alpha : [0, |\alpha| + n] \rightarrow \Delta_d$ by adding α_i steps of integer length i at $\{x = 0\}$ from $\gamma_\alpha(0) = p$ to $\gamma_\alpha(|\alpha|) = (0, I\beta)$. Then we compute $\mu_+(\gamma_\alpha)$ as in definition 5.39 and set $\mu_{\alpha,+}(\gamma) := \frac{1}{I\alpha} \cdot \mu_+(\gamma_\alpha)$.

9.4 Remark

Note that definition 9.3 seems to depend on the order in which we add the α_i steps of lengths i to the path γ to obtain the path γ_α . It will follow however from the alternative description of the positive multiplicity in proposition 9.10 (2) that this is not the case.

We can now define the analogues of the numbers $N_{\text{trop}}^{\alpha,\beta}(d, g)$ in the lattice path setting.


9.5 Definition

Let $d \geq 0$ and g be integers, and let α and β be sequences with $I\alpha + I\beta = d$. We define $N_{\text{path}}^{\alpha,\beta}(d, g)$ to be the *number of λ -increasing paths* $\gamma : [0, 2d + g + |\beta| - 1] \rightarrow \Delta_d$ that start at $(0, d - I\alpha) = (0, I\beta)$ and end at $(d, 0)$, where each such path is counted with multiplicity $\mu_{\alpha,\beta}(\gamma) := \mu_{\alpha,+}(\gamma) \cdot \mu_{\beta,-}(\gamma)$.

Note that as expected (that is as in the complex case) we always have $N_{\text{path}}(d, g) = N_{\text{path}}^{(0),(d)}(d, g)$ by definition.

9.6 Example

The following picture shows that $N_{\text{path}}^{(0,1),(1)}(3, 0) = 4 + 2 + 1 + 1 + 2 = 10$:



9.2. THE CAPORASO-HARRIS ALGORITHM FOR GENERALIZED LATTICE PATHS

Our next aim is to reprove the Caporaso-Harris formula for the numbers $N_{\text{path}}^{\alpha,\beta}(d, g)$ of definition 9.5. The idea of this proof is to count the possibilities of the first step of each path γ , and to multiply this with the number of possibilities how the path can go on. If the first step of the path ends in the second column of Δ_d (see definition 3.69), we want to understand the end of the path as a new path in the smaller triangle Δ_{d-1} . For this, we first need to express the negative and positive multiplicities of a generalized lattice path in a different, non-recursive way. We start with an easy preliminary lemma:

9.7 Lemma


Let α and β be two sequences with $I\alpha + I\beta = d$, and let γ be a generalized lattice path as in section 9.1. If γ has a step that “moves at least two columns to the right”, that is, that starts on a line $\{x = i\}$ and ends on a line $\{x = j\}$ for some i, j with $j - i \geq 2$ then $\mu_{\beta,-}(\gamma) = \mu_{\alpha,+}(\gamma) = \mu_{\alpha,\beta}(\gamma) = 0$.

Proof:

Let γ be a generalized path with a step that moves at least two columns to the right. Then both paths γ'_{\pm} and γ''_{\pm} of definition 5.39 also contain such a step. Hence the lemma follows by induction, because the only end paths δ_{β} (see definition 9.2) and δ_{+} (see definition 5.39) of the recursion to compute the multiplicity of γ which do not have multiplicity 0 do not contain such a step. □

9.8 Remark

We can therefore conclude that any generalized lattice path with non-zero multiplicity has only two types of steps: some that go down vertically and others that move exactly one column to the right (with a simultaneous move up or down):



9.9 Notation

For a path as in remark 9.8 we fix the following notation for the whole chapter: for the vertical line $\{x = i\}$ in the triangle Δ_d (see definition 3.69) we let $h(i)$ be the highest y -coordinate of a point of γ on this line. By α^i we denote the sequence describing the lengths of the vertical steps of γ on this line. For example, for the path shown above we have $h(0) = 1$, $h(1) = 3$, $h(2) = 2$, $h(3) = 1$ and $\alpha^0 = (0)$, $\alpha^1 = (1, 1)$, $\alpha^2 = (1)$, $\alpha^3 = (0)$.

We are now ready to interpret both the positive and negative multiplicity of a generalized lattice path “column-wise”:

9.10 Proposition

Let α and β be two sequences with $I\alpha + I\beta = d$, and let γ be a generalized lattice path as in definition 9.5.

(1) The negative multiplicity of γ is given by the formula

$$\mu_{\beta,-}(\gamma) = \sum_{(\beta^0, \dots, \beta^d)} \left(\prod_{i=0}^{d-1} I^{\alpha^{i+1} + \beta^{i+1} - \beta^i} \cdot \binom{\alpha^{i+1} + \beta^{i+1}}{\beta^i} \right)$$

where the sum is taken over all $(d + 1)$ -tuples of sequences $(\beta^0, \dots, \beta^d)$ such that $\alpha^0 + \beta^0 = \beta$ and $I\alpha^i + I\beta^i = h(i)$ for all i .

(2) The positive multiplicity of γ is given by the formula

$$\mu_{\alpha,+}(\gamma) = \frac{1}{I\alpha} \cdot \sum_{(\beta^0, \dots, \beta^d)} \left(\prod_{i=0}^{d-1} I^{\alpha^i + \beta^i - \beta^{i+1}} \cdot \binom{\alpha^i + \beta^i}{\beta^{i+1}} \right)$$

where the sum is taken over all $(d + 1)$ -tuples of sequences $(\beta^0, \dots, \beta^d)$ such that $\beta^0 = \alpha$ and $d - i - I\beta^i = h(i)$ for all i .

9.11 Remark


Before proving this proposition we would like to interpret its statement geometrically.

The formula of proposition 9.10 (1) counts the number of ways to arrange triangles and parallelograms in Δ_d below γ such that

- the subdivision contains all vertical lines $\{x = i\}$ below γ ; and
- each triangle in the subdivision “is pointing to the left”, that is the vertex opposite to its vertical edge lies to the left of this edge,


where each such subdivision is counted with a multiplicity equal to the product of the double areas of its triangles. We will call such a subdivision a *column-wise Newton subdivision*.

The sequences β^i describe the lengths of the vertical edges in the subdivisions below γ . The binomial coefficients $\binom{\alpha^{i+1} + \beta^{i+1}}{\beta^i}$ in the formula count the number of ways to arrange the parallelograms and triangles, and the factors $I^{\alpha^{i+1} + \beta^{i+1} - \beta^i}$ are the double areas of the triangles. As an example let us consider the path of remark 9.8 with $\beta = (1)$. In this case there is only one possibility to fill the area below γ with parallelograms and triangles as above. This is shown in the following picture:



(corresponding to $\beta^0 = \beta^2 = \beta^3 = (1), \beta^1 = (0)$). As there is one triangle in this subdivision with double area 2 we see that $\mu_{\beta,-} = 2$.

Note that the original definition of the negative multiplicity of a path (see 5.39) also counts Newton subdivisions together with their multiplicity. We called those the *possible Newton subdivisions* for γ in remark 5.40. But the possible Newton subdivisions do not have to be column-wise. In fact, if we reconsider the path from above and compute the possible Newton subdivisions below it, we will also get one of multiplicity 2, but it is not column-wise:



We can see that the column-wise subdivision of the path γ from above *does not* correspond to an actual tropical curve through the special point configuration \mathcal{P}_λ (see section 5.4), whereas the subdivision here does. Therefore we can interpret proposition 9.10 as the combinatorial statement that the number of column-wise subdivisions as described above is equal to the number of possible Newton subdivision for a path γ (counted with multiplicities), and with that equal to the number of tropical curves through \mathcal{P}_λ .

Note that for the positive multiplicity there is in fact no such difference between the possible and the column-wise Newton subdivisions: it can be checked that the possible Newton subdivisions for a path γ *above* γ all contain the vertical lines $\{x = i\}$ above γ and are in fact column-wise.

What is important about proposition 9.10 is that in the column-wise subdivisions we can split off the first column to obtain a similar subdivision of Δ_{d-1} . This will be the key ingredient in the proof of the Caporaso-Harris formula in the lattice path set-up in theorem 9.13.

Proof of proposition 9.10:

We start with part (1). The proof is an induction on the recursive definition of $\mu_{\beta,-}$ in definition 9.2. It is obvious that the end paths in this recursion (the paths that go from $(0, I\beta)$ to $(d, 0)$ along the border of Δ_d) satisfy the stated formula: all these paths have $\beta^0 = (0)$, so the condition $\alpha^0 + \beta^0 = \beta$ requires $\alpha^0 = \beta$. But then the path is one of the paths δ_β of definition 9.2.


Let us now assume that $\gamma : [0, n] \rightarrow \Delta_d$ is an arbitrary generalized lattice path. By induction we can assume that the paths γ' and γ'' of definition 5.39 satisfy the formula of the proposition. Recall that if $k \in [1, n - 1]$ is the first vertex at which γ makes a right

turn then γ' and γ'' are defined by cutting this vertex $\gamma(k)$ (respectively completing it to a parallelogram). By lemma 9.7 we know that $\gamma(k-1)$ (respectively $\gamma(k+1)$) can be at most one column to the left (respectively right) of $\gamma(k)$. But $\gamma(k-1)$ cannot be in the same column as $\gamma(k)$ as otherwise the λ -increasing path γ could not make a right turn at $\gamma(k)$. Therefore $\gamma(k-1)$ is precisely one column left of $\gamma(k)$. There are two possibilities where $\gamma(k+1)$ can lie:

- $\gamma(k+1)$ can be in the same column i as $\gamma(k)$; or
- $\gamma(k+1)$ can be one column right of $\gamma(k)$.

We will prove the statement for these two cases separately.

Starting with the first case, assume that $\gamma(k+1)$ is in column i as $\gamma(k)$. That is, locally the paths γ , γ' and γ'' look as in the following picture:



Then the path γ' has the same values of $h(j)$ and α^j (see notation 9.9) as γ , except for $h(i)$ being replaced by $h(i) - s$ and α^i by $\alpha^i - e_s$, where s is the length of the vertical step from $\gamma(k)$ to $\gamma(k+1)$. The path γ'' has the same values of $h(j)$ and α^j as γ except for $h(i)$ being replaced by $h(i) - s$, α^i by $\alpha^i - e_s$, and α^{i-1} by $\alpha^{i-1} + e_s$.

Using the formula of the proposition for γ' and γ'' (which holds by induction) we compute

$$\begin{aligned} \mu_{\beta,-}(\gamma') &= \sum_{(\beta^0, \dots, \beta^d)} \left(\prod_{j=0}^{d-1} I^{\alpha^{j+1} + \beta^{j+1} - \beta^j - \delta_{i,j+1}e_s} \cdot \binom{\alpha^{j+1} + \beta^{j+1} - \delta_{i,j+1}e_s}{\beta^j} \right) \\ &= \frac{1}{s} \cdot \sum_{(\beta^0, \dots, \beta^d)} \left(\prod_{j=0}^{d-1} I^{\alpha^{j+1} + \beta^{j+1} - \beta^j} \cdot \binom{\alpha^{j+1} + \beta^{j+1} - \delta_{i,j+1}e_s}{\beta^j} \right) \end{aligned}$$

(where the sum is taken over the same β as in the proposition); and

$$\mu_{\beta,-}(\gamma'') = \sum_{(\beta^0, \dots, \beta^d)} \left(\prod_{j=0}^{d-1} I^{\alpha^{j+1} + \beta^{j+1} - \beta^j} \cdot \binom{\alpha^{j+1} + \beta^{j+1} + \delta_{i-1,j+1}e_s - \delta_{i,j+1}e_s}{\beta^j} \right)$$

where the conditions on the summation variables β^i are $\alpha^0 + \delta_{i-1,0}e_s + \beta^0 = \beta$ and $I(\alpha^j - \delta_{i,j}e_s + \delta_{i-1,j}e_s) + I\beta^j = h(j) - s\delta_{i,j}$. We can make these conditions the same as in the proposition by replacing the summation variables β^{i-1} by $\beta^{i-1} - e_s$, arriving at the formula

$$\mu_{\beta,-}(\gamma'') = \sum_{(\beta^0, \dots, \beta^d)} \left(\prod_{j=0}^{d-1} I^{\alpha^{j+1} + \beta^{j+1} - \beta^j} \cdot \binom{\alpha^{j+1} + \beta^{j+1} - \delta_{i,j+1}e_s}{\beta^j - \delta_{i,j+1}e_s} \right).$$

Plugging these expressions into the defining formula

$$\mu_{\beta,-}(\gamma) = s \cdot \mu_{\beta,-}(\gamma') + \mu_{\beta,-}(\gamma'')$$

we arrive at the formula of the proposition (where we use the standard binomial identity $\binom{n-1}{k} + \binom{n-1}{k-1} = \binom{n}{k}$).

For the second case, assume $\gamma(k+1)$ is one column right of $\gamma(k)$. The idea for this case is the same as for the previous case. But the computation is simpler, because the path γ' does not give a contribution due to lemma 9.7. First note that neither the step from $\gamma(k-1)$ to $\gamma(k)$ nor the step from $\gamma(k)$ to $\gamma(k+1)$ can have a negative slope. This is true because otherwise the paths γ'' would also have this negative slope, as we complete the corner to a parallelogram. But the end paths δ_β do not have a step of negative slope, so the claim follows by induction. That is, the heights $h(i)$ of γ in the three columns of $\gamma(k-1)$, $\gamma(k)$ and $\gamma(k+1)$ increase.

Hence locally at k , the paths γ , γ' and γ'' look as in the following picture:



The path γ'' has the same values of $h(j)$ and α^j as γ , except for $h(i)$ being replaced by $h(i) - s + t$, where i is the column of $\gamma(k)$, and s and t are the vertical lengths of the steps before and after $\gamma(k)$. (By the above we know that $s, t \geq 0$.) So using the formula of the proposition again for γ'' (which holds by induction) we have

$$\mu_{\beta,-}(\gamma'') = \sum_{(\beta^0, \dots, \beta^d)} \left(\prod_{j=0}^{d-1} I^{\alpha^{j+1} + \beta^{j+1} - \beta^j} \cdot \binom{\alpha^{j+1} + \beta^{j+1}}{\beta^j} \right)$$

where the conditions on the β^j are $\alpha^0 + \beta^0 = \beta$ and $I\alpha^j + I\beta^j = h(j) - (s - t)\delta_{i,j}$. Note that $s - t > 0$ since γ makes a right turn. As in the previous case we can make the conditions on the β^j the same as in the proposition by replacing the summation variables β^i by $\beta^i - e_{s-t}$. Then we have

$$\mu_{\beta,-}(\gamma'') = \sum_{(\beta^0, \dots, \beta^d)} \left(\prod_{j=0}^{d-1} I^{\alpha^{j+1} + \beta^{j+1} - \beta^j} \cdot \binom{\alpha^{j+1} + \beta^{j+1} - \delta_{i,j+1}e_{s-t}}{\beta^j - \delta_{i,j}e_{s-t}} \right).$$

This is already the formula of the proposition except for the factors

$$\binom{\alpha^i + \beta^i}{\beta^{i-1}} \binom{\alpha^{i+1} + \beta^{i+1}}{\beta^i} \quad \text{being replaced by} \quad \binom{\alpha^i + \beta^i - e_{s-t}}{\beta^{i-1}} \binom{\alpha^{i+1} + \beta^{i+1}}{\beta^i - e_{s-t}}$$


But these terms are the same by the identity $\binom{n+k+l}{n+k} \binom{n+k}{n} = \binom{n+k+l}{n+l} \binom{n+l}{n}$ (note that $\alpha^i = (0)$ for our path γ). This completes the proof.

The proof of case (2) works analogously. However, note that a step of the second type (that is, which passes over two columns) can never be the first left turn, because the path

starts at p in the first column. This is the reason why the subdivisions above γ are indeed column-wise (see remark 9.11). \square

9.12 Remark

Note that it is important for the second step of the proof above that the two boundary lines of Δ_d below and above γ — the line $\{y = 0\}$ respectively the diagonal line from $(0, d)$ to $(d, 0)$ — are indeed straight lines. We cannot generalize the proof to polygons which contain a vertex above respectively below γ , because then the heights of the three columns of $\gamma(k - 1)$, $\gamma(k)$ and $\gamma(k + 1)$ cannot be described as $h(i)$, $h(i) + s$ and $h(i) + s + t$. So we cannot use the identity $\binom{n+k+l}{n+k} \binom{n+k}{n} = \binom{n+k+l}{n+l} \binom{n+l}{n}$. The picture below shows a polygon for which the formula of proposition 9.10 does not hold. The column-wise subdivisions would predict 0 as negative multiplicity for the path. However, the path γ'' is a valid end path, so we get 1.



The proof can be generalized to polygons where the boundaries above and below γ are straight lines, for example to rectangles.

We are now ready to prove the Caporaso-Harris formula in the lattice path setting:

9.13 Theorem

The numbers $N_{\text{path}}^{\alpha, \beta}(d, g)$ satisfy the Caporaso-Harris formula (see definition 3.37).

Proof:

The idea of the proof is to list the possibilities of the first step of the path γ . Let γ be a λ -increasing path from $(0, I\beta)$ to $q = (d, 0)$. As we have seen in remark 9.8 (using lemma 9.7), there are two cases for the first step of γ (corresponding to the two types of summands in the Caporaso-Harris formula):

- The point $\gamma(1)$ is on the line $\{x = 0\}$.
- The point $\gamma(1)$ is on the line $\{x = 1\}$.

In the first case the point $\gamma(1)$ must be $(0, I\beta - k)$ for some $\beta_k \neq 0$ as otherwise the multiplicity $\mu_{\beta, -}(\gamma)$ would be 0. It follows that the restriction $\gamma|_{[1, 2d+g+|\beta|-1]}$ is a path from $(0, d - I(\alpha + e_k))$ and with $\mu_{\alpha+e_k, \beta-e_k}(\gamma|_{[1, 2d+g+|\beta|-1]}) = k \cdot \mu_{\alpha, \beta}(\gamma)$. Therefore the paths γ with $\gamma(1) \in \{x = 0\}$ contribute

$$\sum_{k: \beta_k > 0} k \cdot N_{\text{path}}^{\alpha+e_k, \beta-e_k}(d, g)$$

to the number $N_{\text{path}}^{\alpha, \beta}(d, g)$.

In the second case, we can use proposition 9.10 to compute both the negative and the positive multiplicity as a product of a factor coming from the first column and the (negative respectively positive) multiplicity of the restricted path $\tilde{\gamma} := \gamma|_{[1, 2d+g+|\beta|-1]}$. More

precisely, we have

$$\begin{aligned}\mu_{\alpha,\beta}(\gamma) &= \mu_{\beta,-}(\gamma) \cdot \mu_{\alpha,+}(\gamma) \\ &= \sum_{\beta'} I^{\beta'-\beta} \binom{\beta'}{\beta} \mu_{\beta',-}(\tilde{\gamma}) \cdot \sum_{\alpha'} \binom{\alpha}{\alpha'} \mu_{\alpha',+}(\tilde{\gamma}) \\ &= \sum_{\alpha',\beta'} I^{\beta'-\beta} \binom{\beta'}{\beta} \binom{\alpha}{\alpha'} \cdot \mu_{\alpha',\beta'}(\tilde{\gamma}).\end{aligned}$$

So the contribution of the paths with $\gamma(1) \notin \{x=0\}$ to $N_{\text{path}}^{\alpha,\beta}(d, g)$ is

$$\sum I^{\beta'-\beta} \binom{\beta'}{\beta} \binom{\alpha}{\alpha'} \cdot N_{\text{path}}^{\alpha',\beta'}(d-1, g')$$

where the sum is taken over all possible α', β' and g' . Let us figure out what these possible values are. It is clear that $\alpha' \leq \alpha$ and $\beta \leq \beta'$. Also, $I\alpha' + I\beta' = d-1$ must be fulfilled. As $\tilde{\gamma}$ has one step less than γ we know that $2d+g+|\beta|-1-1 = 2(d-1)+g'+|\beta'|-1$ and hence $g-g' = |\beta'-\beta|-1$. A path $\epsilon: [0, n] \rightarrow \Delta$ from $(0, I\beta)$ to q that meets all lattice points of Δ has $|\beta| + d(d+1)/2$ steps. As γ has $2d+g-1+|\beta|$ steps, $|\beta| + d(d+1)/2 - (2d+g-1+|\beta|) = (d-1)(d-2)/2 - g$ lattice points are missed by γ . But $\tilde{\gamma}$ cannot miss more points, therefore $(d-2)(d-3)/2 - g' \leq (d-1)(d-2)/2 - g$, that is $d-2 \geq g-g'$. \square

9.14 Remark

The same argument can also be applied to other polygons Δ . For example, the analogous recursion formula for $\mathbb{P}^1 \times \mathbb{P}^1$, that is for a $d' \times d$ rectangle $\Delta_{(d',d)}$ reads

$$\begin{aligned}N_{\text{path}}^{\alpha,\beta}((d', d), g) &= \sum_{k:\beta_k>0} k \cdot N_{\text{path}}^{\alpha+e_k, \beta-e_k}((d', d), g) \\ &\quad + \sum I^{\beta'-\beta} \binom{\alpha}{\alpha'} \binom{\beta'}{\beta} N_{\text{path}}^{\alpha',\beta'}((d'-1, d), g')\end{aligned}$$

for all α, β with $I\alpha + I\beta = d$, where the second sum is taken over all α', β', g' such that $\alpha \leq \alpha'$, $\beta' \geq \beta$, $I\alpha' + I\beta' = d$, $g-g' \leq d-1$ and $|\beta-\beta'| = g'-g-1$.

9.3. THE CORRESPONDENCE BETWEEN TROPICAL CURVES WITH ENDS OF HIGHER WEIGHT AND GENERALIZED LATTICE PATHS

As we have already seen in the previous section, the numbers $N_{\text{path}}^{\alpha,\beta}(d, g)$ satisfy the Caporaso-Harris formula. By induction we can conclude the $N_{\text{path}}^{\alpha,\beta}(d, g) = N_{\text{trop}}^{\alpha,\beta}(d, g) = N_{\text{cplx}}^{\alpha,\beta}(d, g)$, because both the numbers $N_{\text{trop}}^{\alpha,\beta}(d, g)$ and $N_{\text{cplx}}^{\alpha,\beta}(d, g)$ satisfy the same recursion formula by 3.38 and 8.6. However, here we prefer to give a direct proof of the statement $N_{\text{trop}}^{\alpha,\beta}(d, g) = N_{\text{path}}^{\alpha,\beta}(d, g)$, which gives more intuition. The idea of the proof is analogous to Mikhalkin's proof of theorem 5.44 stating that $N_{\text{trop}}(d, g) = N_{\text{path}}(d, g)$ (see section 5.4).

9.15 Theorem

For all d, g, α, β we have $N_{\text{trop}}^{\alpha,\beta}(d, g) = N_{\text{path}}^{\alpha,\beta}(d, g)$, where $N_{\text{trop}}^{\alpha,\beta}(d, g)$ is defined in 8.2 and $N_{\text{path}}^{\alpha,\beta}(d, g)$ in 9.5.

Proof:

As usual we choose $\lambda(x, y) = x - \varepsilon y$. Similar to definition 5.46, let \mathcal{P}_λ be a set of $2d + g + |\beta| - 1$ points in restricted general position on the line H orthogonal to the kernel of λ , such that the distance between p_i and p_{i+1} is much bigger than the distance of p_{i-1} and p_i for all i , and such that all points lie below the fixed ends. In other words, if the fixed ends have the y -coordinates $y_1, \dots, y_{|\alpha|}$ then the y -coordinates of p_i are chosen to be less than all $y_1, \dots, y_{|\alpha|}$. Our aim is to show that the number of tropical curves through this special configuration is equal to the number of lattice paths as in section 5.4. Let C be an (unparametrized) tropical curve with the right properties through this set of points. Mark the points where H intersects the fixed ends. Then lemma 5.48 tells us that the edges in Δ dual to the edges of the curve where they meet \mathcal{P}_λ and the new marked points form a λ -increasing path from $p = (0, d)$ to $q = (d, 0)$. The fact that the fixed ends lie above all other points tells us that the path starts with α_i steps of lengths i . So we can cut the first part and get a path from $(0, I\beta)$ to q with the right properties.

Next, let a path $\gamma : [0, 2d + g + |\beta| - 1] \rightarrow \Delta_d$ be given that starts at $(0, d - I\alpha)$ and ends at q . Extend γ to a path $\gamma_\alpha : [0, |\alpha| + n] \rightarrow \Delta_d$ by adding α_i steps of integer length i at $\{x = 0\}$ from $\gamma_\alpha(0) = p$ to $\gamma_\alpha(|\alpha|) = (0, I\beta)$. Add the steps of integer lengths i in an order corresponding to the order of the fixed ends. As in the proof of theorem 5.44, the recursive definition of $\mu_{\beta, -}(\gamma_\alpha)$ corresponds to counting the possibilities for a dual tropical curve in the half plane below H through \mathcal{P}_λ . Passing from γ_α to γ'_α and γ''_α corresponds to counting the possibilities in a strip between H and a parallel shift of H . We end up with a path δ_- which begins with α_i steps of length i and continues with β_i steps of lengths i . This shows that all possible dual curves have the right horizontal ends. Furthermore, $\mu_{\beta, -}(\gamma_\alpha)$ coincides with the number of possible combinatorial types of the curve in the half plane below H times the multiplicity of the part of the curve in the half plane below H . With the same arguments we get that $\mu_{\alpha, +}(\gamma_\alpha)$ is equal to the number of possibilities for the combinatorial type times the multiplicity in the upper half plane. Altogether, we have

$$\begin{aligned}
N_{\text{path}}^{\alpha, \beta}(d, g) &= \sum_{\gamma} \text{mult}_{\alpha, \beta}(\gamma) \\
&= \frac{1}{I^\alpha} \sum_{\gamma} \text{mult}_{\beta, -}(\gamma_\alpha) \cdot \text{mult}_{\alpha, +}(\gamma_\alpha) \\
&= \frac{1}{I^\alpha} \sum_C \text{mult}(C) = \sum_C \text{mult}_{\alpha, \beta}(C) \\
&= N_{\text{trop}}^{\alpha, \beta}(d, g),
\end{aligned}$$

where C runs over all tropical curves with the right properties and γ runs over all paths with the right properties. \square

REFERENCES

1. Kai Behrend, *Gromov-Witten invariants in algebraic geometry*, Invent. Math. **127** (1997), no. 3, 601–617.
2. Kai Behrend and Barbara Fantechi, *The intrinsic normal cone*, Invent. Math. **128** (1997), 45–88.
3. Kai Behrend and Yuri Manin, *Stacks of stable maps and Gromov-Witten invariants*, Duke Math. J. **85** (1996), no. 1, 1–60.
4. Lucia Caporaso and Joe Harris, *Counting plane curves of any genus*, Invent. math. **131** (1998), 345–392.
5. David Cox and Sheldon Katz, *Mirror symmetry and algebraic geometry*, vol. 68, AMS, 1999.
6. Manfred Einsiedler, Mikhail Kapranov, and Douglas Lind, *Non-archimedean amoebas and tropical varieties*, Preprint, math.AG/0408311, 2004.
7. Barbara Fantechi, *Stacks for everybody*, European Congress of Mathematics, 2000, pp. 349–359.
8. Mikael Forsberg, Mikael Passare, and August Tsikh, *Laurent determinants and arrangements of hyperplane amoebas*, Adv. in Math. **151** (2000), 45–70.
9. William Fulton, *Intersection theory*, Ergebnisse der Mathematik und ihrer Grenzgebiete, no. 2, Springer, 1998.
10. William Fulton and Rahul Pandharipande, *Notes on stable maps and quantum cohomology*, Algebraic Geometry, Santa Cruz 1995 (János Kollár et al., ed.), Proceedings of Symposia in Pure Mathematics, no. 62,2, Amer. Math. Soc., 1997, pp. 45–96.
11. Andreas Gathmann, *Absolute and relative Gromov-Witten invariants of very ample hypersurfaces*, Duke Mathematical Journal **115** (2002), no. 2, 171–203.
12. Andreas Gathmann and Hannah Markwig, *The Caporaso-Harris formula and plane relative Gromov-Witten invariants in tropical geometry*, Preprint, math.AG/0504392, 2005.
13. ———, *Kontsevich’s formula and the WDVV equations in tropical geometry*, Preprint, math.AG/0509628, 2005.
14. ———, *The numbers of tropical plane curves through points in general position*, Journal für die reine und angewandte Mathematik (to appear), Preprint, math.AG/0504390, 2005.
15. Ilia Itenberg, Eugenio Shustin, and Viatcheslav Kharlamov, *Welschinger invariant and enumeration of real plane rational curves*, Preprint, math.AG/0303378.
16. Zur Izhakian, *Duality of tropical curves*, Preprint math.AG/0503691.
17. Finn Knudsen, *Projectivity of the moduli space of stable curves, II*, Math. Scand. **52** (1983), 1225–1265.
18. Joachim Kock and Israel Vainsencher, *An invitation to quantum cohomology: Kontsevich’s formula for rational plane curves*, Birkhäuser, 2006.
19. Maxim Kontsevich and Yuri Manin, *Gromov-Witten classes, quantum cohomology and enumerative geometry*, Comm. Math. Phys. (1994), 525–562.
20. Jacques Martinez, *Perfect lattices in euclidean spaces*, Springer, 2003.

21. Grigory Mikhalkin, *Tropical curves and their jacobians*, Preprint available at the homepage <http://www.math.toronto.edu/~mikha/draft.ps>.
22. ———, *Tropical geometry*, Book in progress, available at the homepage <http://www.math.toronto.edu/~mikha/book.ps>.
23. ———, *Enumerative tropical geometry in \mathbb{R}^2* , J. Amer. Math. Soc. **18** (2005), 313–377, Preprint, math.AG/0312530.
24. Takeo Nishinou and Bernd Siebert, *Toric degenerations of toric varieties and tropical curves*, Preprint, math.AG/0409060, 2004.
25. Juergen Richter-Gebert, Bernd Sturmfels, and Thorsten Theobald, *First steps in tropical geometry*, Idempotent Mathematics and Mathematical Physics, Proceedings Vienna (2003).
26. Eugenii Shustin, *A tropical approach to enumerative geometry*, Algebra Anal. **17** (2005), no. 2, 170–214, Preprint, math.AG/0211278.
27. David Speyer, *Tropical geometry*, Ph.D. thesis, University of California, Berkeley, 2005.
28. David Speyer and Bernd Sturmfels, *The tropical grassmannian*, Adv. Geom. **4** (2004), 389–411.
29. ———, *Tropical mathematics*, Preprint, math.CO/0408099, 2004.
30. Luis Felipe Tabera, *Tropical Cramer's rule and tropical Pappus' theorem*, Preprint, math.AG/0409126.
31. Tammo tom Dieck, *Topologie*, de Gruyter, 1991.
32. Ravi Vakil, *Counting curves on rational surfaces*, manuscripta math. (2000).
33. Magnus Dehli Vigeland, *The group law on a tropical elliptic curve*, Preprint, math.AG/0411485, 2004.

INDEX

- $N_{\text{path}}(\Delta, g)$, 107
- $\mathcal{M}_{\text{trop}, 0,4}$, 50, 137, 138, 141, 148, 149
- (α, β) -multiplicity, 152
- 3-valent, 48, 52, 56, 59, 64, 67, 70, 72, 76, 80, 92, 93, 102, 123, 131, 156, 158, 162
- 4-valent, 50, 60, 65, 75, 85, 86, 98, 136
- J_t -holomorphic, 116, 117, 121, 127, 152, 160, 161, 163
- $N_{\text{cplx}}^{\text{irr}}(\Delta, g)$, 46
- $N_{\text{cplx}}^{\text{irr}}(d, g)$, 23
- $N_{\text{trop}}^{\text{irr}}(\Delta, g)$, 70, 71
- $N_{\text{cplx}}(\Delta, g)$, 46
- $N_{\text{cplx}}(d, g)$, 24
- $N_{\text{cplx}}^{\alpha, \beta}(d, g)$, 35
- $N_{\text{cplx}}^{\text{irr}, (\alpha, \beta)}(d, g)$, 42
- $N_{\text{path}}^{\alpha, \beta}(d, g)$, 166
- $N_{\text{trop}}(\Delta, g)$, 89, 105
- $N_{\text{trop}}^{\alpha, \beta}(d, g)$, 152
- $N_{\text{trop}}^{\text{irr}, (\alpha, \beta)}(d, g)$, 157
- Log, 7
- Val, 9
- ev-general position, 75
- λ -increasing path, 105
- $\mathcal{M}'_{\text{trop}, g, n}(\Delta)$, 53, 89
- $\mathcal{M}_{\text{trop}, g, n}$, 49
- $\mathcal{M}_{\text{trop}, g, n}(\Delta)$, 53
- $\mathcal{M}_{\text{trop}, g, n}^{\alpha}(\Delta)$, 54
- $\mathcal{M}_{g, n}$, 18
- $\mathcal{M}_{g, n}(X, \beta)$, 21
- $\overline{\mathcal{M}}_{g, n}^i(\mathbb{P}^2, d)$, 36
- $\overline{\mathcal{M}}_{g, n}$, 20
- $\overline{\mathcal{M}}_{g, n}(X, \beta)$, 21
- π -general position, 136
- π -multiplicity, 135, 137, 138, 142, 143, 145
- $\widetilde{\mathcal{M}}_{\text{trop}, g, n}(\Delta)$, 60
- f -general position, 72
- n -marked abstract tropical curve, 49
- n -marked curve, 18
- n -marked parametrized tropical curve, 51
- n -marked stable map, 21
- abstract curve, 18
- abstract tropical curve, 49
- affine variety associated to a cone, 44
- algorithm of Caporaso and Harris, 36
 - for irreducible curves, 41
 - for tropical irreducible curves, 158
 - lattice path setting, 172
 - tropical, 154
- amoeba, 7, 115
- balancing condition, 16, 51
- boundary of a combinatorial type, 58
- bounded edge, 48
- canonical parametrization of an edge, 50
- Caporaso-Harris formula, 35, 154, 172
- codimension of a combinatorial type, 60, 61, 66, 70, 72, 75, 78, 136, 139
- column-wise Newton subdivision, 168
- combinatorial type, 54, 56–58, 60, 61, 68, 71, 75, 77, 102, 130, 134
- completion of the field of Puiseux series, 9, 122
- complex tropical curve, 121
- component
 - of an unparametrized tropical curve, 93
- cone, 43
- connected
 - graph, 48
 - stable map, 24
- Correspondence Theorem, 114
 - generalized, 159
- degree
 - of a J_t -holomorphic curve, 116
 - of a complex tropical curve, 121
 - of a morphism of polyhedral complexes, 72
 - of a Newton subdivision, 96
 - of a parametrized tropical curve, 52
 - of a projective curve, 23
 - of an unparametrized tropical curve, 93
- dimension
 - expected, 56
 - of a polyhedral complex, 71
 - of a stratum of the moduli space of relevant parametrized tropical curves, 61
 - of the moduli space of quasi stable maps, 37
 - of the moduli space of stable curves, 20
 - of the moduli space of stable maps, 22
- divisors
 - on the moduli space of stable maps, 28
 - on toric surfaces, 46, 52
- dual
 - of a tropical curve, 14, 91, 94
 - of marked points, 101
- dual cone, 44
- dual Newton subdivision, 14, 91, 94
- edge, 48, 96, 101
 - of a graph, 48
 - of a marked unparametrized tropical curve, 101

- of a Newton subdivision, 94
 - of an unparametrized tropical curve, 96
- edge multiplicity, 122, 162
- end, 48
- evaluation map, 22
 - tropical, 66
- exceptional type, 59, 61, 75, 76
- expected dimension, 56

- fan, 43
- flag, 48
- flat loop, 59, 65, 76
- forgetful map, 27
 - tropical, 133

- general position, 17, 71, 102
 - ev-general position, 75
 - π -general position, 136
 - f -general position, 72
 - in a polyhedral complex, 71
 - restricted general position, 102, 114, 150
 - tropical, 70
- generalized Correspondence Theorem, 159
- generalized lattice path, 167
- genus
 - of a J_t -holomorphic curve, 116
 - of a complex tropical curve, 121
 - of a graph, 48
 - of a Newton subdivision, 98
 - of an unparametrized tropical curve, 93
- graph, 48
 - 3-valent, 48
 - connected, 48
 - genus, 48
- Gromov-Witten invariants, 23

- Hausdorff metric, 119

- integer length, 14, 96, 97, 102, 165, 166
- intersection multiplicity, 143, 151
- irreducible
 - Newton subdivision, 99
 - unparametrized tropical curve, 93

- Kapranov's Theorem, 13, 91
- Kontsevich's formula, 30
 - tropically, 129

- lattice, 55
- lattice path, 105
 - generalized, 167
- length of an edge, 50
- marked curve, 18

- marked edges, 101
- marked point, 52
- marked unparametrized tropical curve, 100
 - edge, 101
 - vertex, 101
- Mikhalkin's Correspondence Theorem, 71, 114
- moduli space
 - of (connected) parametrized tropical curves, 53, 57
 - of (not necessarily connected) parametrized tropical curves, 89
 - of abstract tropical curves, 49
 - of quasi-stable maps, 36
 - of rational tropical curves, 131
 - of relevant parametrized tropical curves, 60
 - of stable curves, 20
 - of stable maps, 21
- morphism of polyhedral complexes, 72
- multiplicity, 60, 68, 71, 73, 80, 116, 131, 135, 145
 - π -multiplicity, 135, 137, 138, 142, 143, 145
 - edge multiplicity, 122, 162
 - intersection multiplicity, 143
 - of a generalized lattice path, 166
 - of a lattice path, 105
 - of a morphism of polyhedral complexes at a point, 72
 - of a parametrized tropical curve, 68, 80, 99
 - of a simple subdivision, 99
 - of a vertex, 67, 99
 - of an unparametrized tropical curve, 93

- Newton polygon, 13, 96
- Newton subdivision, 14, 95
 - column-wise, 168, 169
 - degree, 96
 - edge, 94
 - genus, 98
 - irreducible, 99
 - possible for a path, 106, 169
 - regular, 95
 - simple, 95
 - vertex, 94
- non-archimedean norm, 9

- parametrized tropical curve, 51
 - degree, 52
 - regular, 56
 - relevant, 60
 - simple, 91
 - superabundant, 56
- polyhedral complex, 71
- possible Newton subdivisions for γ , 106, 169

- pre-stable n -marked curve, 19
- pre-stable n -marked map, 20
- primitive integral vector, 7, 13, 16, 52, 60, 92, 94, 102, 104, 122, 153, 161
- Puiseux series, 9, 122
- quasi-stable map, 36
- reducible
 - rational tropical curve, 141
 - unparametrized tropical curve, 93
- regular
 - Newton subdivision, 95
 - parametrized tropical curve, 56
- relative Gromov-Witten invariants, 37
 - tropical, 152
- relevant, 60, 61, 66, 68, 71, 72, 76, 86, 92, 102, 104, 130, 131
- relevant subset, 60
- relevant type, 60
- restricted general position, 102
- rigid, 67, 70
- root vertex, 55
- Severi variety, 37
- simple
 - marked parametrized tropical curve, 100
 - Newton subdivision, 95
 - parametrized tropical curve, 91
- Speyer's Theorem, 56, 91
- spine, 117
- stability, 19
- stabilization
 - of a curve, 25
 - of a map, 26
- stable curve, 20
- stable map, 21
- stack, 20, 22
- string, 67, 139, 154
- superabundant, 56, 59
- toric surface, 43, 114
 - associated to a fan, 44
- torus, 43
- tropical algorithm of Caporaso and Harris, 154
 - for irreducible curves, 158
 - lattice path setting, 172
- tropical curve
 - associated to a curve over the field of Puiseux series, 9, 91
 - associated to a tropical polynomial, 11, 91
 - complex tropical curve, 121
 - parametrized, 51
 - reducible, 141
 - unparametrized, 92
- tropical enumerative problem, 70
- tropical evaluation map, 66
- tropical forgetful map, 133
- tropical general position, 70
- tropical Kontsevich's formula, 129
- tropical polynomial, 10, 91
- tropical relative Gromov-Witten invariants, 152
- tropical semiring, 10
- tropicalization, 117
 - of a polynomial, 11
- unbounded edge, 48
- unparametrized tropical curve, 92
 - component, 93
 - degree, 93
 - edge, 96
 - genus, 93
 - irreducible, 93
 - marked, 100
 - multiplicity, 93
 - reducible, 93
 - vertex, 96
- valence of a vertex, 48
- valuation, 9
- vertex, 13, 14, 48, 94, 96, 101
 - of a marked unparametrized tropical curve, 101
 - of a Newton subdivision, 94
 - of an unparametrized tropical curve, 96
- Viro's Patchworking method, 124
- virtual fundamental class, 22
- weight, 13, 15, 16, 52, 92, 96, 97, 102, 152, 165

CURRICULUM VITAE

Name	Hannah Markwig
Date of birth	November 19th, 1980
Address	home: Am Berg 5, 67705 Stelzenberg, Germany office: FB Mathematik, TU Kaiserslautern, 67653 Kaiserslautern, Germany email: markwig@mathematik.uni-kl.de
1999	Abitur at the Peter-Petersen-Gymnasium (secondary school) in Mannheim, Germany
1999	TU Kaiserslautern undergraduate scholarship
1999-2003	Studienstiftung des deutschen Volkes (German National Academic Foundation)
1999-2001	Study of Mathematics with a minor in Physics at the TU Kaiserslautern, Germany
2001	Vordiplom at the TU Kaiserslautern
2002	Otto-Wipprecht-Foundation study abroad grant
2002	Semester abroad at the University of California, Berkeley, USA
2001-2003	Study of Mathematics with a minor in Philosophy at the TU Kaiserslautern, Germany
2003	Diploma (German equivalent of a master's degree) in mathematics at the TU Kaiserslautern
since 2003	PhD studies at the TU Kaiserslautern, advisor Andreas Gathmann
2004	The Rheinland-Pfalz State Graduate Scholarship
2004-2006	Research assistant at the TU Kaiserslautern, DFG (Deutsche Forschungsgesellschaft - German Research Organization) research fellowship
since 2006	Scientific assistant at the TU Kaiserslautern