Tropical Curves

1 Introduction

These notes are based are an introduction to the study of tropical curves. The approach is informal, with emphasis on intuition and examples rather than complete proofs. Our primary written reference is [**?**]. As always comments, corrections and suggestions are very welcome.

2 Amoebas of Curves

The basic idea of tropical geometry is to study a complex plane curve by looking at its image in $\mathbb{R}^2$ under the map:

$$\begin{aligned} \varphi : (\mathbb{C}^*)^2 &\to \mathbb{R}^2 \\ (x, y) &\mapsto (\log|x|, \log|y|) \end{aligned}$$

If $C \subseteq (\mathbb{C}^*)^2$ is a complex curve then its image $\varphi(C)$ is called the **amoeba** of C. Since C has two real dimensions, we might expect the same to be true for its amoeba. Usually this will be the case, so that the amoeba of a complex curve forms a (real) surface in $\mathbb{R}^2$.

Remark 2.1. This fact is not completely obvious and actually does not hold in all situations. For instance, it is clear that φ contracts any radial circle $S^1_x \subseteq \mathbb{C}^*_x \subseteq (\mathbb{C}^*_{xy})^2$ to a point, and similarly for any circle $S^1_y \subseteq (\mathbb{C}^*_{xy})^2$. Thus there are entire tori $S^1_x \times S^1_y$ which are mapped to a single point under φ, so that φ does not always preserve dimensions. However, if the curve we choose is sufficiently generic then it will be transverse to these tori, and then the image will have two real dimensions as desired.

Aside 2.2. The map φ can be viewed as the moment map for the tautological action of the Clifford torus $S^1 \times S^1$ on $(\mathbb{C}^*)^2$, where $(\mathbb{C}^*)^2$ is equipped with the symplectic form $\omega = \Re(\mathrm{d}x/x \wedge \mathrm{d}y/y)$.

We will now look at a couple of examples. In what follows we will use uppercase X and Y to denote the (real) co-ordinates on the codomain.

Example 2.3. Our first example is extremely simple: letting $C = \{y = c\}$ in $(\mathbb{C}^*)^2$ for some constant c, we have:

$$\varphi(C) = \{Y = \log|c|\}$$

Thus the amoeba associated to the line C is simply another line (though this time in $\mathbb{R}^2$). This has real dimension 1, whereas C has real dimension 2; the reason for this discrepancy is that C contains radial circles S^1_x which are contracted down to a point by φ (see Remark 2.1 above). In fact, $\varphi : C \to \varphi(C)$ is a fibration with circle fibres of the form S^1_x.

Example 2.4. Moving away from this degenerate example, let us consider a generic line $C = \{ax + by = c\} \subseteq (\mathbb{C}^*)^2$ (we assume from now on that a, b, c are positive real numbers; the reason for this will soon become clear). As a variety this is isomorphic to $\mathbb{P}^1$ minus three points (a line in $\mathbb{C}^2$ is $\mathbb{P}^1$ minus a single point, and we lose two more points by excluding the cases $x = 0$ and $y = 0$).

In order to study $\varphi(C)$ we examine what happens when $|x|$ or $|y|$ tends to 0 or ∞; note that this is equivalent to $\log |x|$ or $\log |y|$ tending to $-\infty$ or ∞ respectively, so that we are really examining the asymptotics of the amoeba. There are essentially three cases to consider.

First, consider the case $|x| \gg 0$ (equivalently $|y| \gg 0$), so that $|a||x| \simeq |b||y|$. Projecting along φ, this is equivalent to $A + X \simeq B + Y$ (where of course $A = \log a, B = \log b$). Thus as x and y approach ∞, the amoeba of C approaches the line $\{X + A = Y + B\}$ in $\mathbb{R}^2$.

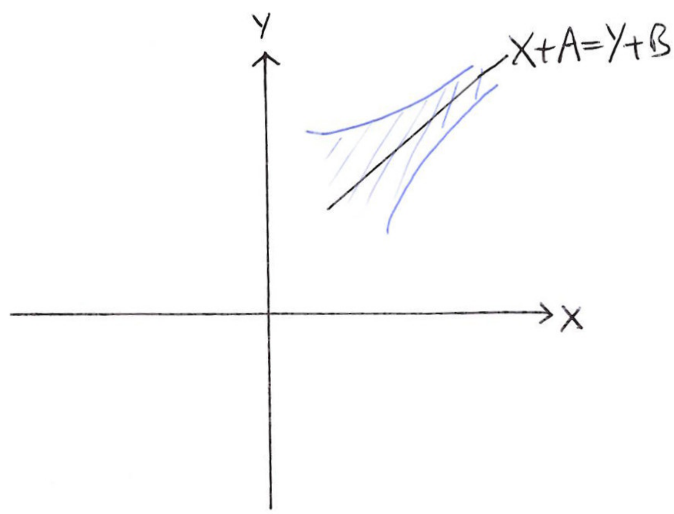

The second case to consider is when $|x| \ll 1$, so that $|b||y| \simeq |c|$. Again, projecting along φ we see that this is equivalent to $B + Y \simeq C$, so that as x approaches 0 (that is, as X approaches $-\infty$) the amoeba of C approaches the line $\{Y + B = C\}$:

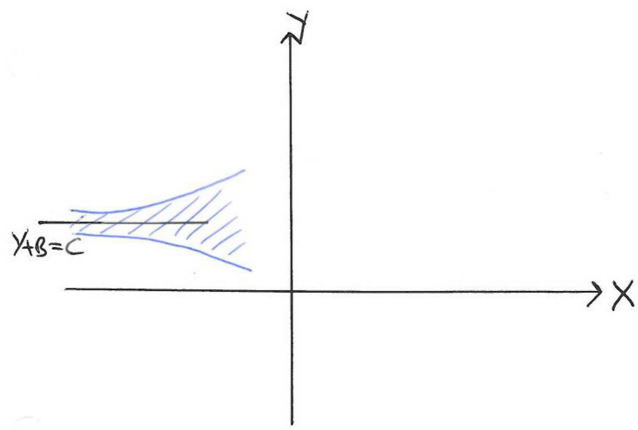

The final case to consider is when $|y| \ll 1$. The same arguments as in the second case apply, and so we see that as y approaches 0 (that is, as Y approaches $-\infty$) the amoeba of C approaches the line $\{X + A = C\}$.

Putting all of these together, we see that our amoeba looks something like this:

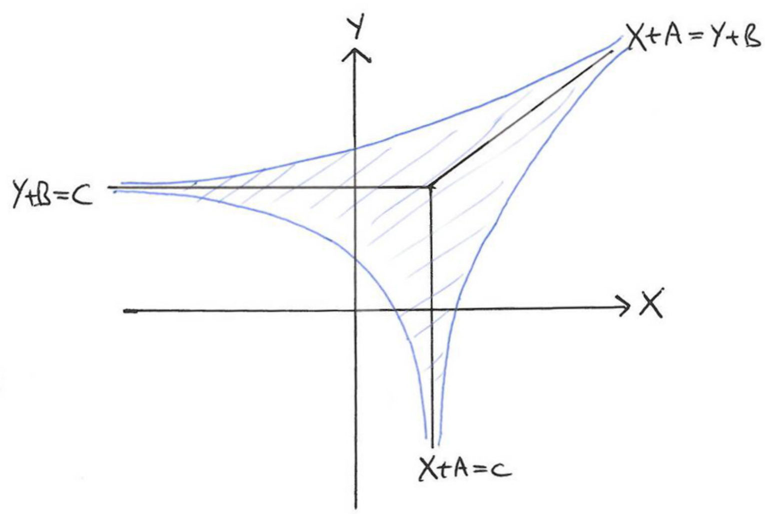

Notice the key role of the lines $\{X + A = Y + B\}$,$\{Y + B = C\}$ and $\{X + A = C\}$, and more precisely the subsegments of these lines drawn in the above figure (which indicate the direction in which they govern the amoeba's asymptotics). Ideally we would like to forget about the amoeba itself and just concentrate on these line segments. This is made precise in the notion of the graph of an amoeba, which we now turn to.

3 The Graph of an Amoeba

The idea is to somehow flow the amoeba down to its underlying "graph" (an example of which is illustrated in the above figure). To achieve this, we modify our map φ by taking logarithms with arbitrary bases

$$\begin{aligned} \varphi_t : (\mathbb{C}^*)^2 &\to \mathbb{R}^2 \\ (x, y) &\mapsto (\log_t |x|, \log_t |y|) = \left(\frac{\log |x|}{\log t}, \frac{\log |y|}{\log t} \right) \end{aligned}$$

and then letting $t \to \infty$. As t increases this forces the central "belly" of the amoeba into a smaller and smaller region, whereas the thin "tentacles" take up a greater proportion of the space. In the limit we should get the graph of the amoeba.

However there is a snag: if the central "vertex" of our amoeba is not at the origin (consider for example the previous figure, where the vertex is at $(C - A, C - B)$) then this limiting process will end up moving the vertex to the origin, since all the scaling happens relative to the co-ordinate axes. To get around this we replace our curve C by a family of curves C_t. In the example of a generic line considered above, this family is given by:

$$C_t = \{t^A x + t^B y + t^C = 0\} \subseteq (\mathbb{C}^*)^2$$

We then take the limit of the amoebas $\varphi_t(C_t)$ as $t \to \infty$ and this gives us the graph as desired. This will be called the **tropical curve** associated to (C_t), or the **tropicalisation** of (C_t). We think of this as some sort of degeneration of the original curves.

Thus, the natural objects to tropicalise are not individual curves but rather families of curves (C_t) depending on a parameter $t \in \mathbb{R}$. However, we cannot tropicalise every such family: we require that the equations for C_t only involve (possibly fractional) powers of t. (The reason for this will become clearer when we consider tropical curves in the context of Puiseux series: see Section 9.)

Despite the definition being in terms of families, we will often want to tropicalise a single curve C (as in the above example of a generic line). To do this, we proceed as we did in this example, by first passing to the naturally associated family (C_t) and then taking the tropicalisation of this family. We call the resulting tropical curve the **tropicalisation** of C.

Computing tropicalisations using the current definition is no easy task. It would be good if we could find a more explicit expression for the tropicalisation of a curve C, which doesn't require us to construct the family (C_t). This leads us to the notion of a tropical polynomial and its corner locus.

4 Corner Loci and Tropical Curves

Let us concentrate on the example of a generic line. Our curves C_t are defined by the polynomials:

$$P_t = t^A x + t^B y + t^C$$

We'll start by examining what happens to $\log_t |P_t|$ as $t \to \infty$ (later we will use this to get an explicit formula for the tropicalisation). As in our earlier discussion, we consider several different cases depending on the values of x and y.

Let us begin by considering the case $|x| \gg |y| \gg -\infty$. Then $\log_t |P_t| \simeq \log_t |t^A x| = A + X$, and this approximation gets more and more accurate as t increases. In the limit we have:

$$\lim_{t\to\infty} \log_t |P_t| = A + X$$

Similarly there is the case $|y| \gg |x| \gg -\infty$. An identical argument shows that in this region we have:

$$\lim_{t\to\infty} \log_t |P_t| = B + Y$$

Finally we consider the case where both $|x|$ and $|y|$ are small. Then $\log_t |P_t| \simeq \log_t |t^C| = C$, and again this approximation becomes an equality in the limit:

$$\lim_{t\to\infty} \log_t |P_t| = C$$

Note that in each of these cases, it was the biggest of $A + X$, $B + Y$ and C which dominated and ended up appearing in the limit. Put differently:

$$\lim_{t\to\infty} \log_t |P_t| = \max\{A + X, B + Y, C\}$$

Denoting this quantity by $P = P(X, Y)$ (and viewing it as a function on $\mathbb{R}^2$ rather than on $(\mathbb{C}^*)^2$) we have a piecewise-affine function whose different values partition the plane as in the figure below:

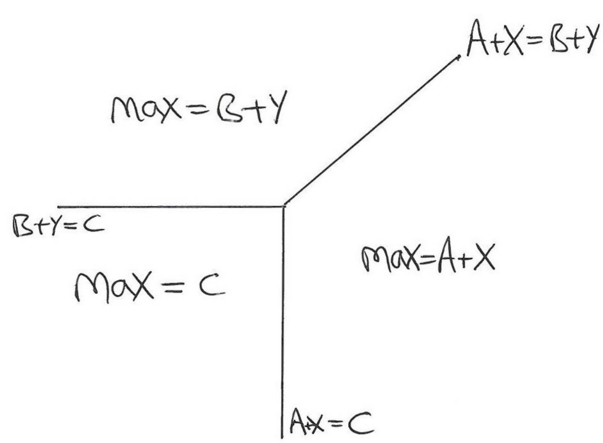

(Note that these regions in $\mathbb{R}^2$ correspond to the regions in $(\mathbb{C}^*)^2$ which we considered in the case analysis above.)

Having done all this, we see that the tropical curve (that is, the graph of the amoeba) has appeared as the "walls" of the partition. More precisely, it is the **corner locus** for P: the locus on which P is not differentiable. Put differently, it is the set of points at which the maximum is attained by at least two of the inputs simultaneously.

What is the point of all this? Notice that the function P can be obtained directly from the data of

$$C = \{ax + by + c = 0\} = \{e^A x + e^B y + e^C = 0\}$$

without having to consider the associated family of curves. Thus we can find the tropicalisation of C without having to worry about any families.

5 The Tropical Semiring

To recap: passing from a complex plane curve to the associated tropical curve involves transforming $ax + by + c$ into $\max\{A + X, B + Y, C\}$. That is, we have replaced $+$ by max and $\times$ by $+$.

This leads us immediately to the definition of the **tropical semiring**. As a set this is just $\mathbb{T} = \mathbb{R}$, but equipped with the operations $\max\{- , -\}$ for addition and $+$ for multiplication. It satisfies all the usual axioms for a ring, except that there is no additive identity element. For more on tropical algebra, see [**?**].

Aside 5.1. Some authors also include the element $-\infty$ in $\mathbb{T}$. In this case there is an additive identity, but there are no additive inverses so we still do not have a proper ring.

Aside 5.2. There is of course some scope for confusion here, since the symbol "+" could be understood either as addition in the tropical semiring or as ordinary addition (which is multiplication in the tropical semiring). We'll be somewhat loose with our notation, relying mostly on context. Thus for instance we might write a tropical polynomial in $\mathbb{T}[X, Y]$ in two different but equivalent ways:

$$AX^2 + BXY + CX + D = \max\{A + 2X, B + X + Y, C + X, D\}$$

With the discussion of the previous section in mind, we define a **tropical curve** to be the corner locus (in $\mathbb{R}^2$) of some tropical polynomial. Given a plane curve $C \subseteq (\mathbb{C}^*)^2$ we can view the defining equation as a tropical polynomial: its corner locus is then the tropicalisation of C.

6 Weights and the Balancing Condition

Consider the (non-generic) plane conic $C = \{ax^2 + by^2 + c = 0\} \subseteq (\mathbb{C}^*)^2$. The tropicalisation is then the corner locus of $\max\{2X + A, 2Y + B, C\}$:

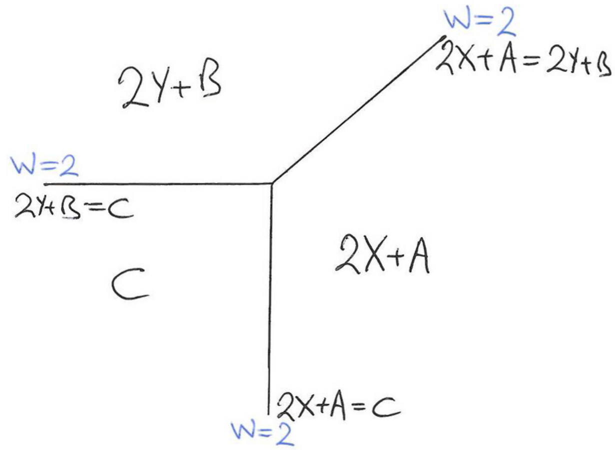

Note that this is (up to a possible translation) the same set of points as in the example of a generic line. However, the projections from the plane curves onto their tropicalisations are different: in the case of the line, the projection is a simple S^1-fibration, whereas in the case of the conic it is the double cover of an S^1-fibration.

This difference is encoded in the data of "weights" which we attach to each wall of the tropical curve. Thus each wall for the conic has weight 2 (see the above figure) whereas those for the line will have weight 1.

More generally: to each wall which is the locus where $A_{ij} + iX + jY = A_{kl} + kX + lY$, we attach the **weight** $w = \gcd\{k-i, l-j\}$. The idea is that above this wall the projection is a w-fold cover of an S^1-fibration, where the S^1 fibres point in the direction orthogonal to $(i-k, j-l)$.

Topological considerations upstairs in C show that these weights must satisfy a combinatorial condition at each vertex of the tropical curve, called the **balancing condition**. If we have walls $e_1, \ldots, e_r$ meeting at a vertex p, with weights $w_1, \ldots, w_r$, then the balancing condition says that

$$\sum_{i=1}^{r} w_i v_i = 0$$

where each v_i is a primitive lattice vector generating e_i.

Example 6.1. Consider the case of the nongeneric conic seen above. Assume for simplicity that the vertex is at the origin. Then the walls have primitive generators $(-1,0), (0,-1), (1,1)$, and the balancing condition is satisfied:

$$2(-1,0) + 2(0,-1) + 2(1,1) = 0$$

If we attach weights to a tropical curve as above, the balancing condition will always be satisfied. There is a sort of converse: any weighted rational graph in $\mathbb{R}^2$ which satisfies the balancing condition at each vertex is a tropical curve (actually, there is an additional condition relating to the number of "tentacles" which go off in each direction; we won't really go into this, though the discussion in the next section should hint at what that condition might be).

7 Degree and the Number of Tentacles

Example 7.1. As an exercise, compute the tropicalisation of a generic conic C. You should find that, so long as the coefficients of x, y and xy are chosen large enough, the result is:

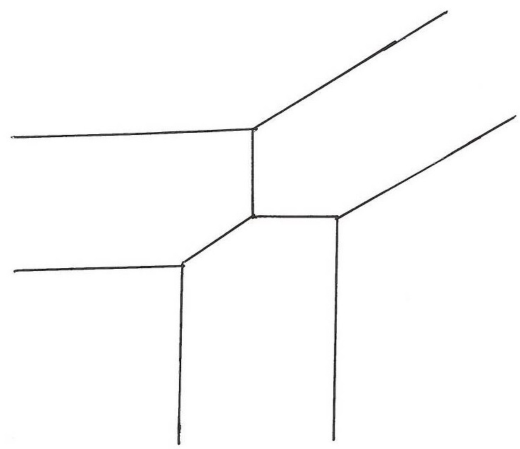

Note that there are two tentacles going off in each direction. It is easy to see why: if we set $x = 0$ we obtain a quadratic equation in y; thus there should be two points in C of the form $(0, y)$. Taking log we see that there are two points in the tropicalisation of the form $(-\infty, Y)$, and these are the two tentacles of the tropicalisation going off to the left. There is a similar story if we set $x = 0$.

More generally, the same argument shows that the tropicalisation of a generic degree d curve should have d tentacles going off in each direction. In fact, this is true even if we consider nongeneric curves, so long as each tentacle is counted with its weight as described in the previous section.

Example 7.2. From the previous discussion, it is plausible (though not completely proven) that a generic cubic will look like:

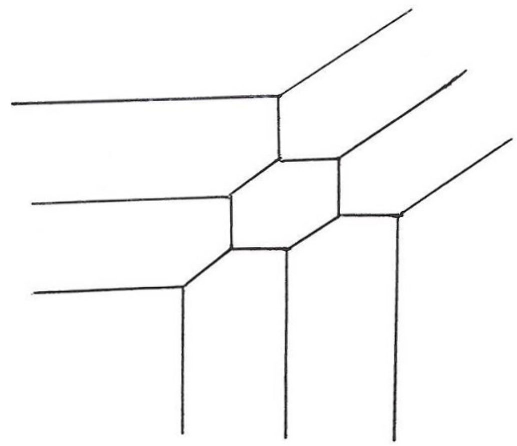

Notice the cycle: thinking of our complex cubic C as an S^1-fibration over this picture, the cycle implies that C will have genus 1. Compare this to the previous example, where the lack of cycles means C has genus 0.

This is what we would expect from the degree-genus formula for plane curves. As an exercise, extend these arguments to arbitrary degree and check that they agree with what we already know from the degree-genus formula (this exercise will become a lot easier once you've read the next section).

8 Newton Polygons

Since tropical curves are purely combinatorial objects, their systematic use can often reduce problems in algebraic geometry to combinatorics. In order to then answer these combinatorial questions, it is sometimes helpful to take a different point of view: that of Newton polygons and their polyhedral decompositions. In a sense this picture is dual to the description of tropical curves as rational graphs.

We begin with a definition. If $f \in k[x, y]$ is a polynomial (over any field or indeed any ring) we define the **Newton polygon** of f to be the convex hull of those multi-indices whose corresponding cofficients are nonzero in f

$$\mathrm{Newt}(f) = \mathrm{conv}\left\{(i, j) \in \mathbb{N}^2 : a_{ij} \neq 0\right\}$$

(where a_{ij} is the coefficient of $x^i y^j$ in f). This is a lattice polygon in $\mathbb{N}^2$. If f is generic of degree d, then $\mathrm{Newt}(f)$ will be a right triangle of side length d.

If $P \in \mathbb{T}[x, y]$ is a tropical polynomial, we define a polyhedral decomposition of $\mathrm{Newt}(P)$ as follows. Each lattice point of $\mathrm{Newt}(P)$ corresponds to a input in $P = \max\{-\}$. We put an edge between two such lattice points if the corresponding inputs are equal at some point (that is, if the corresponding regions in $\mathbb{R}^2$ intersect along a wall of the tropical curve). Similarly, we put in 2-dimensional faces bounding a set of edges if the corresponding walls in the tropical curve meet at a vertex.

This gives us a polyhedral decomposition of $\mathrm{Newt}(P)$. It follows immediately from the construction that the dual polyhedral decomposition is just the tropical curve associated to P.

Example 8.1. Consider the generic line: the Newton polygon is a right triangle of side length 1, so of course the polyhedral decomposition is trivial. Dualising we obtain the tropical curve we have seen before:

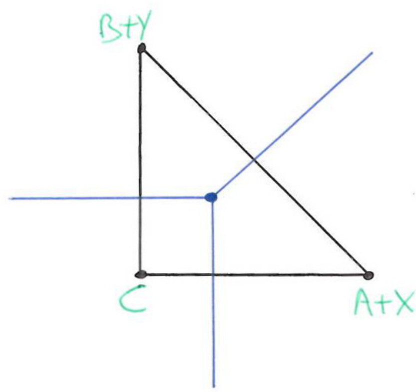

Example 8.2. Consider the generic conic: the Newton polygon is a right triangle of side length 2, and the polyhedral decomposition is the maximally fine triangulation. Dualising we obtain the tropical curve:

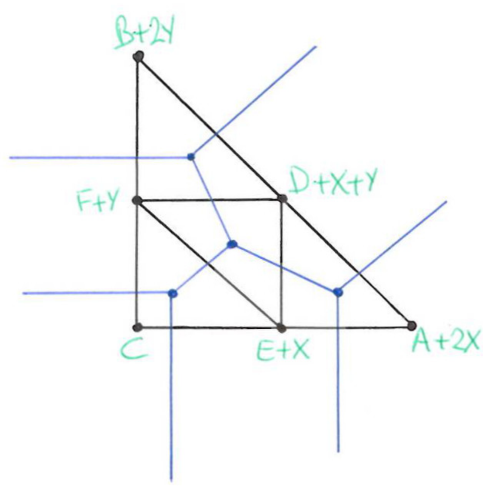

On its own this construction is not terribly useful, since we require full knowledge of the tropical curve in order to determine the polyhedral decomposition of Newt(P). Better would be if we could obtain some alternative description of the polyhedral decomposition, since then we could dualise to obtain the tropical curve, practically for free.

As it turns out there is such a description. We begin by labeling each lattice point of Newt(P) by the coefficient of the corresponding monomial in P. These determine a convex piecewise-linear function on the polygon, by requiring that at each lattice point the value of the function equals the label.

This then determines a polyhedral decomposition of Newt(P): the 2-dimensional faces of the decomposition are precisely the regions on which the function is linear. It can be shown that this decomposition is the same as the one defined previously.

Example 8.3. Consider again the case of a generic conic. We have the labeled Newton polygon

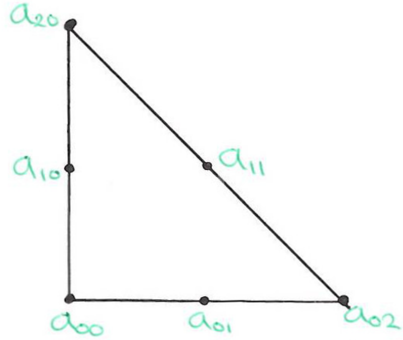

and if we choose our coefficients generic enough, the corresponding function will only be linear on the small subtriangles of the following decomposition:

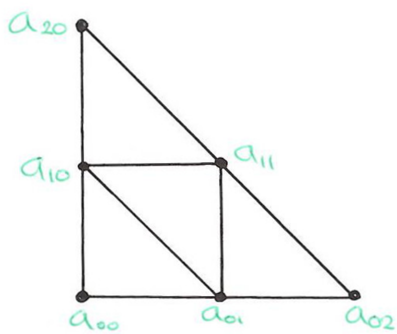

Dualising this we recover the tropical curve of a generic conic as in Example 8.2.

Example 8.4. On the other hand, if we had chosen our coefficients non-generically (for instance if we set $a_{00} = a_{01} = a_{10} = a_{11}$) we could obtain the following:

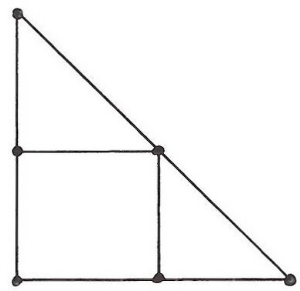

This is obtained from the previous polyhedral decomposition by fusing together two of the sub-polygons; we think of this as some sort of degeneration of the generic case. The tropical curve we get is:

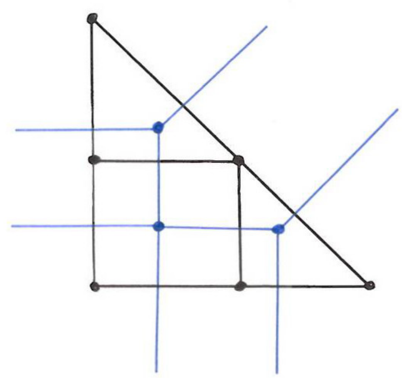

9 Puiseux Series

In Section 3 we introduced in a rather ad hoc way families of curves (C_t), in order to avoid all the vertices of the tropical curve getting sent to the origin. There is a more elegant way of constructing tropicalisations in which the appearance of the family (C_t) is far more natural: namely, we think of this family of curves in $(\mathbb{C}^*)^2$ as being a single curve in K^2, where K is the field of Puiseux series.

The Puiseux series are defined to be formal power series (with coefficients in $\mathbb{C}$) indexed by $\mathbb{Q}$

$$\sum_{q \in \mathbb{Q}} a_q t^q$$

such that the set of $q \in \mathbb{Q}$ with $a_q \neq 0$ is bounded above and has only a finite number of denominators appearing in its elements. This is an algebraically closed field of characteristic zero, and so when we do algebraic geometry over K the resulting theory will be similar to that of the complex numbers.

Aside 9.1. In fact, K is the algebraic closure of the field of Laurent series, where we require all but finitely many of the positive terms to vanish.

Coming back to the family of curves (C_t), we now view the parameter t appearing in the equation for C_t as the formal variable in K. Thus the family (C_t) corresponds to a single curve in $(K^*)^2$.

Under this identification, the limit of the functions $\log_t$ as $t \to \infty$ is given by simply picking off the highest nonzero power of t. We call this val, so that:

$$\begin{aligned} \mathrm{val} : K^* &\to \mathbb{R} \\ \sum_{q \in \mathbb{Q}} a_q t^q &\mapsto \max\{q : a_q \neq 0\} \end{aligned}$$

In analogy with our earlier constructions, we define:

$$\mathrm{Val} = \mathrm{val}^2 : (K^*)^2 \to \mathbb{R}^2$$

Thinking now of (C_t) as a curve in $(K^*)^2$, we see that $\mathrm{Val}(C_t)$ coincides with the tropicalisation which we defined earlier as the limit of $\varphi_t(C_t)$.

Thus we can redefine a tropical curve as being the image under Val of an algebraic curve in $(K^*)^2$.

Aside 9.2. Strictly speaking, we should take the closure in $\mathbb{R}^2$ of the image under Val, since this function takes values in $\mathbb{Q}^2$.