



HAL
open science

On the topology of a real analytic curve in the neighborhood of a singular point

Étienne Ghys, Christopher-Lloyd Simon

► **To cite this version:**

Étienne Ghys, Christopher-Lloyd Simon. On the topology of a real analytic curve in the neighborhood of a singular point. *Asterisque*, 2020, 415, pp.1-33. 10.24033/ast.1097 . hal-03273688

HAL Id: hal-03273688

<https://hal.science/hal-03273688>

Submitted on 30 Jun 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On the topology of a real analytic curve in the neighborhood of a singular point

Étienne GHYS and Christopher-Lloyd SIMON

April 15, 2018

Dedicated to the memory of Jean-Christophe Yoccoz

Statement of the main result

In this paper, we propose a complete description of the topology of *real analytic planar curves* in the neighborhood of a singular point.

Denote by $\mathbb{R}\{x, y\}$ the factorial ring of germs of real analytic functions defined in some neighborhood of $(0, 0) \in \mathbb{R}^2$. The germ of a real analytic planar curve is defined by an equation $F(x, y) = 0$, where $F \in \mathbb{R}\{x, y\}$ vanishes at the origin. If F is an irreducible element in this ring, the topology of the curve \mathcal{C}_F defined by F is well known. Either it only contains the origin (as for $x^2 + y^2 = 0$) or there is a *local homeomorphism* of the plane, mapping \mathcal{C}_F to (the germ of) a straight line (as for instance in $x^3 - y^2 = 0$). In this second case, \mathcal{C}_F intersects small circles centered at the origin in exactly two points.

In general F is a product $F_1^{n_1} \cdots F_k^{n_k}$ of irreducible non-associated factors F_i . Our curve \mathcal{C}_F is therefore the union of the \mathcal{C}_{F_i} 's, which are usually called the *branches* of \mathcal{C}_F . Since we are only interested in the topology of \mathcal{C}_F , we can discard those F_i 's such that \mathcal{C}_{F_i} only contains the origin. Two distinct factors F_i yield two branches which only intersect at the origin. Hence the analytic curve \mathcal{C}_F intersects small circles centered at the origin in an even number of points, grouped in pairs, each pair being associated to a branch. This yields a *chord diagram* which is by definition an even number of distinct points on the circle, grouped in pairs, up to an orientation preserving homeomorphism of the circle. Such a diagram is pictured by a certain number of chords with distinct endpoints in a circle.

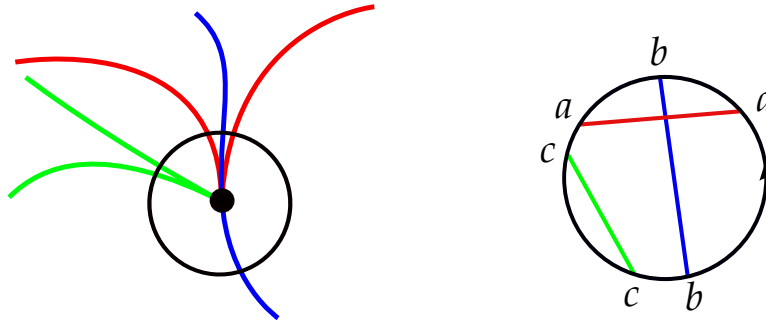


Figure 1: A curve with three branches and its associated chord diagram

The main theorem of this paper characterizes the chord diagrams arising from some analytic curve \mathcal{C}_F .

Theorem. *A chord diagram is associated to some analytic curve if and only if it does not contain one of the “forbidden diagrams” shown in Figure 2 as a sub-chord diagram.*

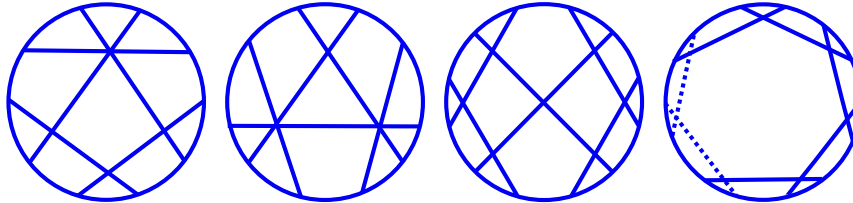


Figure 2: Forbidden diagrams: \diamond , $\diamond\diamond$, \triangle , and C_n ($n \geq 5$)

The genesis of this paper

Consider four distinct polynomials P_1, P_2, P_3, P_4 in $\mathbb{R}[x]$. Order them in such a way that $P_1(x) < P_2(x) < P_3(x) < P_4(x)$ for small *negative* values of x . Then define the permutation π on $\{1, 2, 3, 4\}$ such that $P_{\pi(1)}(x) < P_{\pi(2)}(x) < P_{\pi(3)}(x) < P_{\pi(4)}(x)$ for small *positive* values of x . In 2009, Maxim Kontsevich explained to the first author EG that among the 24 permutations on $\{1, 2, 3, 4\}$ exactly two cannot be obtained by this construction: $(1, 2, 3, 4) \mapsto (2, 4, 1, 3)$ or $(3, 1, 4, 2)$. EG easily generalized this to any number of polynomials and proved that a permutation on $\{1, \dots, n\}$ can be obtained from n polynomials if and only if it does not “contain” one of Kontsevich’s permutations. This was published as an elementary paper [5]. We will give a different proof later in section 2.

It was then very natural to look at the topological configurations of the branches of a real analytic curve in the neighborhood of a singular point. Trying to solve this problem, EG found an explicit algorithm determining if a given chord diagram is *analytic*, i.e. is associated to the branches of some real analytic singular point. In particular, it followed

that the above forbidden chord diagrams were indeed not analytic. One can always delete some branches of an analytic curve, so that a sub-chord diagram of an analytic diagram is of course analytic. In particular, a diagram containing one of the forbidden examples is non-analytic. The question of knowing whether these examples were the only “minimal” forbidden configurations remained open.

Since this proof was enjoyable and involved classical methods, EG decided to write a book proposing a leisurely promenade towards this partial result, intended for undergraduate students. The second author CS was such a student and read a preliminary draft of that book. He proposed to look at the problem from another side, explained below, and this new point of view enabled both authors to complete the proof of the above theorem in a joint effort. Therefore the final version of the book contains an additional chapter, describing this result [6].

The present paper contains two sections. The first section provides an algorithmic description of the analytic chord diagrams and the second uses the first to prove the main result. This paper is very close to the corresponding chapters of the book. We essentially “compressed” these chapters in order to get more efficiently to the main goal.

1 Analytic chord diagrams: an algorithm

In this section, we get an algorithmic description of the analytic chord diagrams, that we defined as those which are determined by the branches of planar real analytic curves.

1.1 Polynomial interchanges: algorithmic description

The only purpose of this subsection is to discuss quickly the much simpler situation of permutations arising from polynomials in $\mathbb{R}[x]$ which were the starting point of this paper. This serves as a motivation and gives a pattern for the general strategy, somewhat different from that in [5].

Let π be a permutation of $\{1, \dots, n\}$ ($n \geq 2$). We say that π is a *polynomial interchange* if there exist n polynomials P_1, \dots, P_n in $\mathbb{R}[x]$ such that

$$P_1(x) < P_2(x) < \dots < P_n(x)$$

for small negative x and

$$P_{\pi(1)}(x) < P_{\pi(2)}(x) < \dots < P_{\pi(n)}(x)$$

for small positive x .

We describe an elementary algorithm that determines if a given permutation is a polynomial interchange. In the next section, we will characterize polynomial interchanges as those permutations which do not contain the two forbidden Kontsevich permutations.

Lemma. *For any polynomial interchange, at least two consecutive integers have consecutive images.*

The proof is easy. Denote by $v(P) \in \mathbb{N} \cup \{\infty\}$ the valuation (at 0) of a polynomial $P \in \mathbb{R}[x]$, i.e. the lowest degree of a non zero monomial in P (and ∞ if $P = 0$). Choose polynomials P_1, \dots, P_n as above. For every integer N , the relation $v(P_i - P_j) \geq N$ is an equivalence relation \mathcal{R}_N on $\{1, \dots, n\}$. Each equivalence class $I \subset \{1, \dots, n\}$ is an interval. Indeed, suppose that $i < j < k$ and that $i, k \in I$. We know that $P_i(x) < P_j(x) < P_k(x)$ for small negative x . It follows that $v(P_j - P_i) \geq v(P_k - P_i) \geq N$ so that $j \in I$. The same argument, for small positive x , implies that $\pi(I)$ is also an interval. Let N_0 be the largest value of N for which equivalence classes of \mathcal{R}_N are not reduced to singletons. Let I be an equivalence class of \mathcal{R}_{N_0} with at least two elements. Since all the valuations $v(P_i - P_j)$ are equal to N_0 for i, j in I , the permutation π is either increasing or decreasing from I to $\pi(I)$, depending on the parity of N_0 . The lemma follows if one chooses two consecutive elements in I . \square

Note in particular that the two permutations $(1, 2, 3, 4) \mapsto (2, 4, 1, 3)$ or $(3, 1, 4, 2)$ are not polynomial interchanges.

Theorem. *The following algorithm decides if a permutation π is a polynomial interchange:*

1. *If no pair of consecutive integers have consecutive images then π is not a polynomial interchange.*
2. *If there is such a pair, merge it to a singleton. This produces a permutation with one object less. Continue.*
3. *If you end up with the trivial permutation on one object, then the original permutation was a polynomial interchange.*

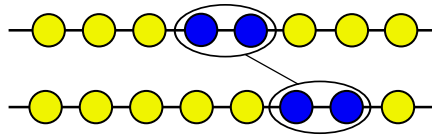


Figure 3: Merging a pair of consecutive elements

If π is a polynomial interchange and if $\{i, i + 1\}, \{\pi(i), \pi(i + 1)\}$ are merged into singletons, we produce a permutation π' on $n - 1$ objects which is obviously another polynomial interchange associated to the polynomials $P_1, \dots, P_i, P_{i+2}, \dots, P_n$. Conversely, given a polynomial interchange associated to $P_1, \dots, P_i, \dots, P_{n-1}$, one can define

$$P'_i(x) = P_i(x) + (-x)^N$$

for a sufficiently large value of N , even or odd, and consider the permutation associated to the n polynomials $P_1, \dots, P_i, P'_i, \dots, P_{n-1}$. This shows that if the merged permutation is a polynomial interchange, then so was the initial permutation. \square

1.2 Chord diagrams

One can think of a chord diagram as a cyclic word of length $2n$ in which every letter occurs exactly twice (the labels of the letters being irrelevant). To be more pedantic (and precise), we are discussing fixed-point free involutions on $\mathbb{Z}/2n\mathbb{Z}$ up to conjugacies by cyclic permutations. We can also draw n chords in a circle.

The total number of *chord diagrams* of length $2n$ has been studied in many papers, like in [12] with strong motivations from knot theory. The problem would be easy if, instead of a cyclic word, we looked for standard (non-cyclic) words of length $2n$ in which every letter occurs exactly twice and whose names are irrelevant. Indeed, write the first letter of the word and then choose any of the remaining $2n - 1$ locations for the other letter which is identical to the first. Then write the second letter in the first available free place and choose the other identical letter in any of the $2n - 3$ remaining locations. Etc. Therefore the total number of these words is $(2n - 1) \cdot (2n - 3) \cdots 3 \cdot 1$. These numbers are sometimes called *double factorials* and denoted by $(2n - 1)!!$.

It would be tempting to divide $(2n - 1)!!$ by $2n$ to take into account the cyclic permutations, but some words admit symmetries and this makes the exact combinatorics more subtle. In any case, it follows from these considerations that the number of chord diagrams of length $2n$ grows super-exponentially in n . We will see that a very tiny proportion of chord diagrams are *analytic*, in the sense that they arise from the singularity of a planar analytic curve.

1.3 A necessary condition

We will say that a chord in a diagram is *solitary* if it connects two consecutive points of the diagram as in the first picture in Figure 4. Two chords are *parallel* (resp. *antiparallel*) if they are as in the second (resp. third) picture, i.e. if the corresponding letters a, b occur in the cyclic word as $\cdots ab \cdots ba \cdots$ (resp. $\cdots ab \cdots ab \cdots$). Finally, two chords as in the fourth picture constitute a *pitchfork* ($\cdots a \cdots bab \cdots$). Letter a is the stick and letter b is the fork. With these notations we can state the fundamental lemma.

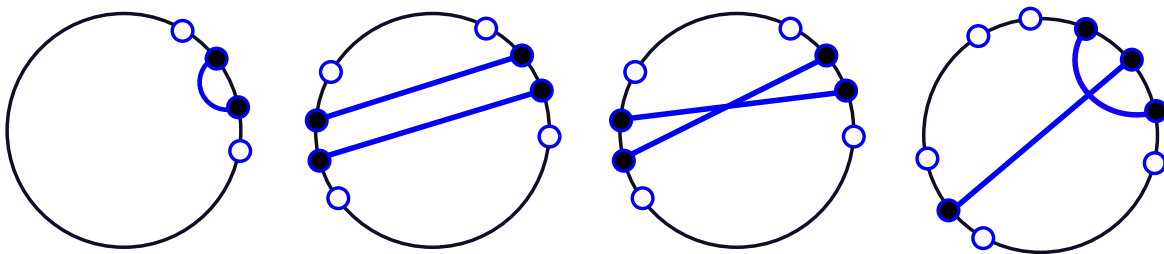


Figure 4: solitary, parallel, antiparallel and pitchfork

Fundamental lemma: *Any analytic chord diagram (with at least two chords) contains a solitary chord, a pair of parallel or an antiparallel chords, or a pitchfork.*

As a simple corollary, there is no singular analytic plane curve whose branches intersect a small circle as in Figure 5.

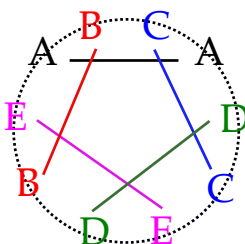


Figure 5: A forbidden diagram with five chords.

Start with some singular point of some analytic curve in the (real) plane. Blow it up a first time¹. The result is a curve in some Moebius band, whose singular points are on the exceptional divisor, core of the band. If things go well, the singular point splits into several singular points, presumably simpler. Let us blow up all of them. It could happen that after one blow up, there is still a unique singular point on the divisor. Then, blow it up a second time. Let us continue the blowing up process as many times as necessary. We know that after some time, the singularity will be *resolved*. This means that the strict transform of the initial curve is now a collection of n disjoint smooth analytic curves intersecting transversally the exceptional divisor.

This exceptional divisor is a union of real projective lines which are circles intersecting transversally. Consider the graph whose vertices are these projective lines and where an edge connects two vertices if the projective lines intersect. The inductive process of desingularization shows that this graph is a tree. Indeed, at each step we blow up a point which can be either a smooth point of the exceptional divisor, or an intersection of two projective lines. In the first case, a new leaf is grafted to a tree and in the second case, an edge is split into two edges. The first projective line, coming from the first blow up, can be chosen as the root of this tree.

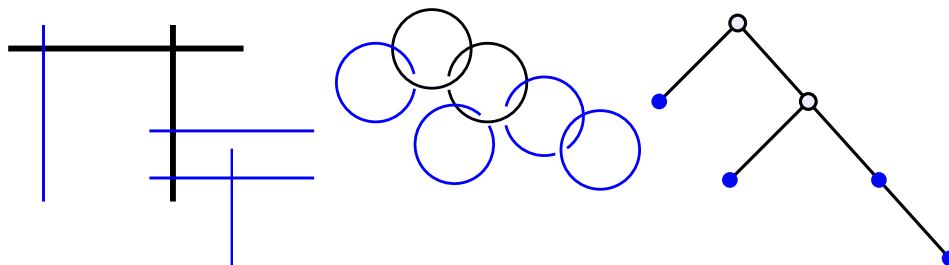


Figure 6: Exceptional divisor made of six projective lines and the associated tree.

¹For the basics about the resolution of real analytic singularities, see the corresponding chapter of [6].

In order to make sure that at the end of the process each projective line contains *at most* one point of the strict transform, it will be convenient to blow up once more each of the n points on the exceptional divisor, if necessary, introducing new projective lines. We can even suppose that the components of the divisor meeting the strict transform are leaves of the desingularization tree.

Let us sum up. Given some analytic curve \mathcal{C} defined in a neighborhood of $(0, 0)$ in \mathbb{R}^2 by some equation $F(x, y) = 0$, we can construct the following objects.

- A surface S with a connected oriented boundary.
- An exceptional divisor $E \subset S$, consisting of a certain number of circles intersecting transversally, each pair meeting at most once. The associated intersection graph is a rooted tree. The embedding $E \subset S$ is a homotopy equivalence.
- The strict transform of \mathcal{C} : a finite disjoint union $\hat{\mathcal{C}}$ of *smooth* analytic arcs β_1, \dots, β_n in S intersecting transversally E . The intersection of $\hat{\mathcal{C}}$ with a component of E is empty if this component is not a leaf of the tree, and contains at most one point if it is a leaf. We can assume moreover that $\hat{\mathcal{C}}$ is transversal to the boundary of S and that each arc β_i intersects the boundary in two points.
- A blowing down analytic map $\Psi : S \rightarrow \mathbb{R}^2$, collapsing E to the origin, which is a diffeomorphism from $S \setminus E$ onto some small punctured disc, and which collapses $\hat{\mathcal{C}}$ to our singular curve \mathcal{C} .

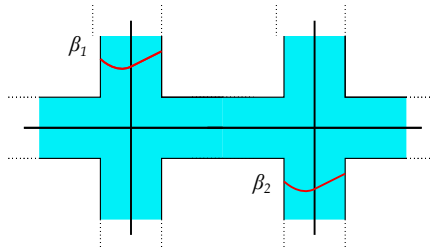


Figure 7: A part of the blown-up surface S .

Each loop in S can be orienting or disorienting. Let γ be a closed immersed curve in a surface, passing once through a point x . When the surface is blown up at x , the self-intersection *modulo 2* of the strict transform of γ (in the blown up surface) is equal to the self-intersection of γ (in the original surface) plus 1. In the inductive construction, when a projective line appears for the first time in the exceptional divisor, it is the core of a Moebius band, of self-intersection 1. Later on, some of its points may be blown up. Each blowing up permutes the orienting/disorienting status of a component in the divisor. Figure 6, is obtained after six blowing ups: six lines, a tree with six vertices, and six circles.

Some components of E intersect the desingularized curve \hat{C} : they define some leaves in the desingularization tree. We say that those leaves are *colored*. Observe that some leaves might not be colored.

Note that if we choose some orientation for each component of E , the corresponding tree is planar so that the children of any node are linearly ordered. Changing the orientation reverses this order.

Look for example at the necklace in Figure 8.

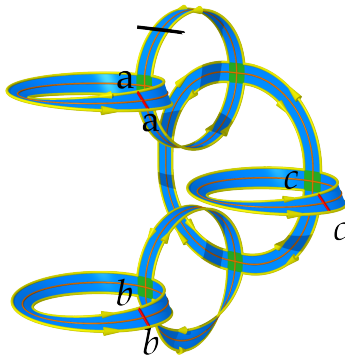


Figure 8: The surface S , neighborhood of the exceptional divisor.

Six blow ups produced six bands, two are orientable and four are not. The exceptional divisor consists of the six cores of the corresponding bands. The desingularized curve is made up from the three red arcs, labeled a, b, c , each intersecting the boundary of S in two points. On top, we see in black the strict transform of the y axis. Going around the boundary of S we can read the corresponding analytic chord diagram. Just follow the arrow and read $abacbc$.

1.4 Proof of the fundamental lemma

First, observe that deleting a chord in an analytic diagram transforms it into some other analytic diagram. This corresponds to deleting a branch.

Now, start with an analytic chord diagram w and consider the desingularization tree of some associated planar singularity. There is a projection ρ of the surface S onto the exceptional divisor E which is a homotopy equivalence. The fiber $\rho^{-1}(x)$ of a point $x \in E$ is an arc connecting two points of the boundary if x is a regular point, and a cross if x is the intersection of two circles.

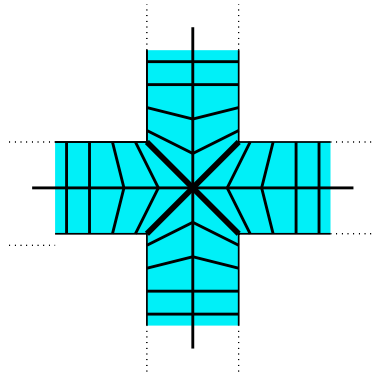


Figure 9: Projection of S onto the exceptional divisor.

Let L be a node of the tree, i.e. one of the projective lines that constitute the exceptional divisor E . There is a unique chain of nodes going from L to the root. Cut two disjoint arcs in S as in Figure 10, in order to disconnect L from the root in S . The four endpoints of these arcs decompose the circle boundary of S into four intervals. Two of them (colored in red) correspond to “what is below L ” in the tree. Going around the boundary of S and reading the chord diagram, we therefore find two disjoint intervals of letters, below L , whose union is stable under the involution sending each occurrence of a letter to the other. Note that these intervals could be empty if there were no colored leaves below L . However, if there is a colored leaf below L , then at least one of the two intervals isn’t empty, though it could happen that only one is not empty.

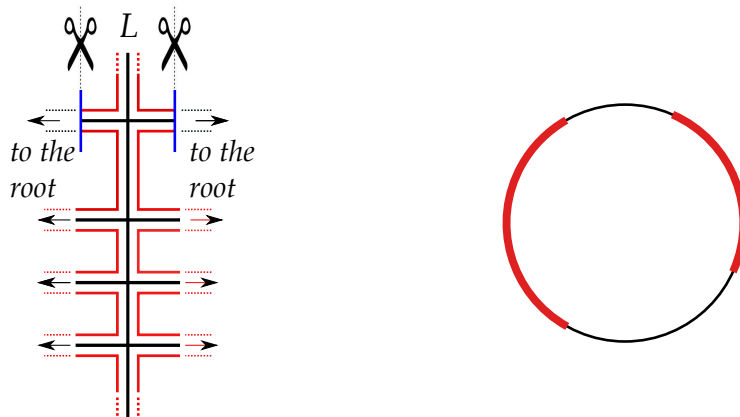


Figure 10: Pruning

In summary, every node L in the tree defines a chord diagram $w(L)$ which is a sub-diagram of the original diagram w and which is “connected” in the sense that its letters form one or two intervals in w .

Think of a rooted tree as a genealogy tree, the root being the founding member of the family. Each node has a certain number of descendants, some of them being colored leaves. Let L be one of the youngest members of the family having at least two colored leaves as descendants. Among the children of L , let L_1, \dots, L_k be the list of those having at least one colored descendant (ordered in this way along L). We have $k \geq 2$ since otherwise one of the children of L would have at least two colored descendants. For the same reason, each L_i has a unique colored descendant.

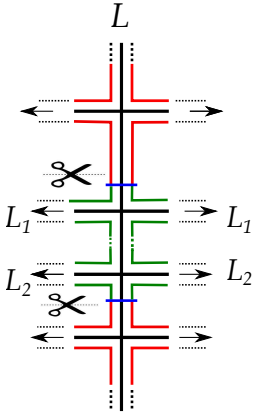


Figure 11: Pruning colored leaves

Now cut the surface S as in Figure 11, disconnecting L_1 and L_2 from the root and from all other colored leaves. As before, this defines two (green) intervals on the boundary of S whose union contains exactly four points of our initial diagram, associated to two chords. Two chords in two intervals can be organized in the following fifteen ways.

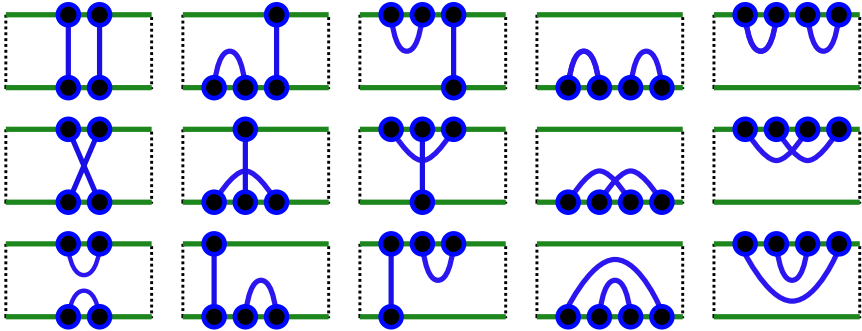


Figure 12: Fifteen cases

In each case, there is a solitary chord, a pitchfork, or a pair of parallel or antiparallel chords. This ends the proof of the fundamental lemma.

□

1.5 More non-analytic diagrams

We have observed that deleting letters in analytic chord diagrams preserves analyticity. A diagram is called *basic non-analytic* if it is not analytic but all of its proper sub chord diagrams are. Clearly a chord diagram is analytic if and only if it does not contain a *basic non-analytic chord diagram*.

Theorem. *There is an infinite number of basic non-analytic chord diagrams.*

Here is an example that will be denoted by C_n ($n \geq 5$). Consider the $2n$ points of $\mathbb{Z}/2n\mathbb{Z}$ naturally ordered on the circle. The chord diagram pairs $2k$ and $2k + 3$ for $k = 1, \dots, n$. For $n = 5$, this is our previous example of non-analytic diagram with five chords. This diagram C_n ($n \geq 5$) is not analytic by the fundamental lemma. We still have to show that if one letter is deleted, the remaining diagram is analytic. For this, we need a *sufficient criterion of analyticity*.

Theorem. *The following algorithm decides if a chord diagram is analytic:*

1. *If there is no solitary chord, no pitchfork, and no pair of parallel or antiparallel chords, the diagram is not analytic.*
2. *If there is a solitary chord, delete it and continue. If there is a pitchfork, delete the fork and keep the stick. If there is a pair of parallel or antiparallel chords, delete one of them and continue.*
3. *If you end up with the one-chord diagram, then the original one was analytic.*

The proof is easy. If w is analytic, \bar{w} is also analytic since it corresponds to deleting a branch. For the converse, we have to show that if w is a diagram and if \bar{w} is the new diagram obtained after one step of the algorithm with one chord less, then w is analytic if \bar{w} is analytic. So we have to *add* an additional branch.

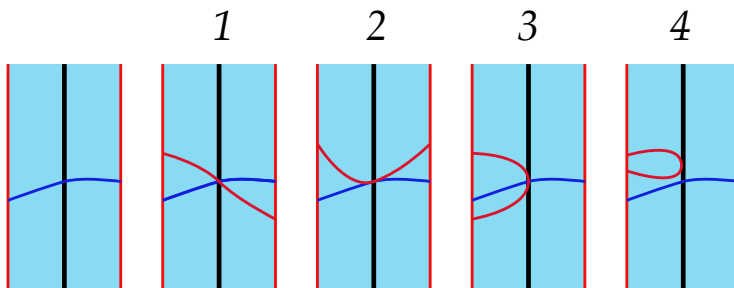


Figure 13: Adding a branch

Choose a desingularization S of a singular point associated to \bar{w} . A chord corresponds to some smooth arc γ connecting two points on the boundary, transverse to the divisor

at a point x . As in 1/, add a new analytic smooth (red) curve γ' in S , very close and transverse to γ as well as to the divisor. The blowing down of S produces a new singular point with one more branch. Clearly, the new diagram has one more chord which is parallel or antiparallel to the initial (blue) chord, depending on the orientations of the boundary. Now choose γ' as in 2/, with a quadratic tangency with γ at x and you get the other parallel or antiparallel situation. To create a pitchfork with a given (blue) handle, just add a smooth curve γ' close to γ with a quadratic tangency with the divisor as in 3/. Finally, if you want to add a solitary chord right after some given letter, proceed as in 4/. \square

Let us use our algorithm to prove that C_n ($n \geq 5$) is *basic non-analytic*. Deleting one chord, we get a non cyclic chain of $n - 1$ chords. The first two chords u, v define a pitchfork, so we may delete the first chord and continue until there is only one chord left. This diagram is therefore analytic. So indeed, the set of basic non-analytic diagrams is infinite.

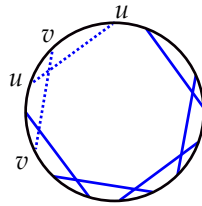


Figure 14: A chain of chords

1.6 With a computer

In order to count analytic chord diagrams, we can use a computer to test small values of n . We start by listing all possible words of length $2n$ in which each letter occurs twice. The only subtlety is to take into account the cyclic character of the word under consideration. Here is the result for $n \leq 7$, in the following table:

n	2	3	4	5	6	7
Words	3	15	105	945	10395	135135
Chord diagrams	2	5	18	105	902	9749
Up to symmetry	2	5	17	79	554	5283

– *Words* means “linear words of length $2n$ in which each letter occurs twice”. There are $(2n - 1)!!$ of those.

– *Chord diagrams*, as we have defined them, are words up to cyclic permutations.

– The image of a chord diagram by a symmetry with respect to some diameter is another diagram, which may be the same diagram or not. The item “up to symmetry” counts the number of cyclic words up to these dihedral symmetries.

We can then count the number of *analytic* diagrams using the algorithm that was described earlier. The result is:

n	2	3	4	5	6	7
Analytic diagrams	2	5	18	102	817	7641
Up to symmetry	2	5	17	76	499	4132

It follows that for $n \leq 4$, all diagrams are analytic.

Among the 105 5-chord diagrams, only the 3 examples in Figure 15 are not analytic.

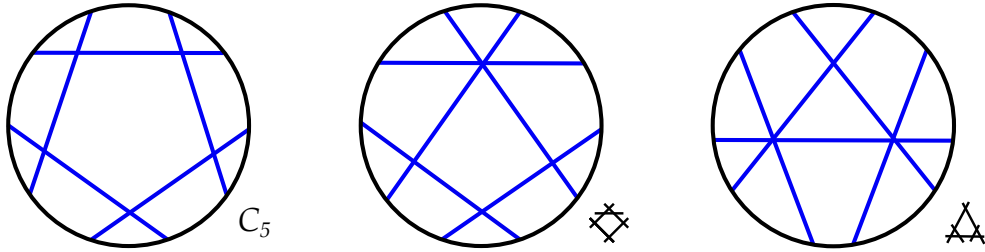


Figure 15: The 3 basic non-analytic chord diagrams with 5 chords.

The first diagram is already familiar, under the name C_5 . We denote the others by \diamond and \triangle . It wasn't difficult to guess the first but the other two were discovered by our computer.

Among the 902 diagrams with 6 chords, 85 are not analytic. However, the non-analyticity of most of them is due to the fact that one of their sub-diagrams is not analytic. Only two 6-chord diagrams are *basic non-analytic*.

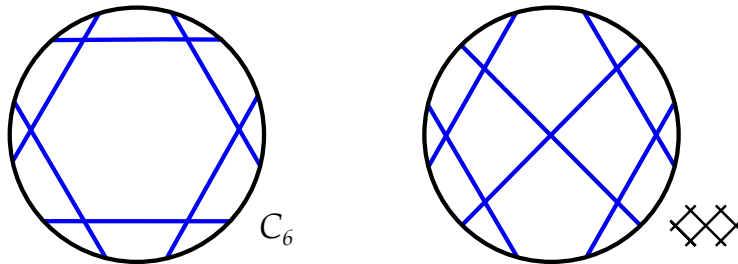


Figure 16: The 2 basic non-analytic chord diagrams with 6 chords.

Observe that the first one is the member C_6 of the infinite family we exhibited. It corresponds to $\mathbb{Z}/12\mathbb{Z}$ where every even number $k \pmod{12}$ is connected to $k + 3 \pmod{12}$. The second will be denoted by $\diamond\diamond$.

Among the 9749 diagrams with 7 chords, 2108 are not analytic. The only basic non-analytic example is C_7 .

Later, we will show that the computer did indeed find *all* basic non-analytic diagrams.

1.7 Marked chord diagrams

It will be convenient to introduce a slight strengthening of the notion of analytic chord diagrams.

When we proved that our algorithm decides if a diagram is analytic, the key point was the possibility of inserting a new branch. It turns out that more complicated singularities can also be inserted, as we now explain. Consider a desingularization of some curve \mathcal{C} as before, so that we have a surface S , a divisor E , and a collection of smooth curves β_1, \dots, β_n intersecting E transversally at p_1, \dots, p_n , where n is the number of real branches. Choose one of these points, say p_1 . Choose now some *other* singular curve \mathcal{C}_1 , with n_1 real branches, and assume that it does not contain the y axis. Delete β_1 and replace it by a copy of \mathcal{C}_1 in the surface S , in such a way that the y axis for \mathcal{C}_1 is mapped into the divisor E and the singular point of \mathcal{C}_1 is mapped to p_1 . We can now blow down the union of this copy of \mathcal{C}_1 and β_2, \dots, β_n . The result is a new singular point, with $n + n_1 - 1$ branches: one of the branches of \mathcal{C} has been replaced by a copy of \mathcal{C}_1 .

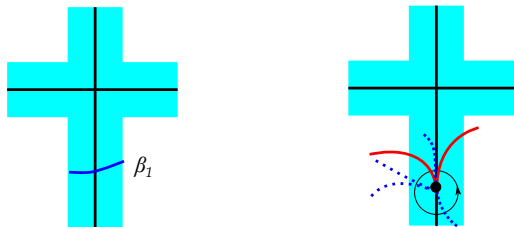


Figure 17: Inserting a singular point.

Let us examine the effect of this operation on the associated chord diagram. Looking at the diagram associated to \mathcal{C}_1 , we see that the y axis decomposes the word of length $2n_1$ in two components, *Left* and *Right*. In the new chord diagram with $2(n + n_1 - 1)$ letters, one pair of identical letters from the old diagram with $2n$ letters has been replaced by two intervals, which are *Left* and *Right*. We should bear in mind that during this process, the orders of the letters in *Left* and *Right* might have been reversed. Indeed, the two intersections of the oriented boundary of S with β_1 might be of different signs. Moreover, the insertion of \mathcal{C}_1 in S can be done in four ways since S is not orientable and E is not oriented.

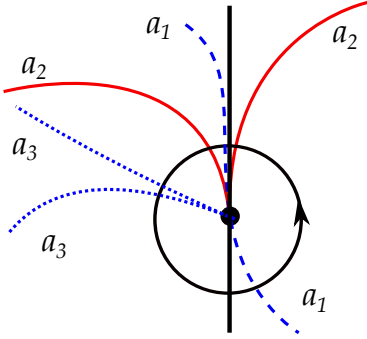


Figure 18: $Right = a_1 a_2$, $Left = a_1 a_2 a_3 a_3$

Note that *Left* or *Right* could be empty. In an equation $F(x, y) = 0$, we can replace (x, y) by $(-x, y)$ or $(x, -y)$ or (x, xy) . The transformation $(x, y) \mapsto (x, xy)$ preserves each vertical line, collapses the axis $x = 0$ to the origin, and reverses the orientation for $x < 0$. This is no surprise: it is a blow down map. The square of this transformation preserves the orientation on each vertical line (for $x \neq 0$).

Of course, we can proceed in the same way with all other branches of \mathcal{C} , using other singular curves $\mathcal{C}_2, \dots, \mathcal{C}_n$.

All these remarks suggest the following definition.

Definition. A marked chord diagram is a collection of $2n$ distinct points $a_1^{\pm 1}, \dots, a_n^{\pm 1}$ in the union of two opposite sides of a square $\{-1, 1\} \times [-1, 1]$ (up to orientation preserving homeomorphisms of each side).

Note the additional features if one compares with standard diagrams. Marked diagrams have a right and a left part. Moreover, each chord a_i is now *labeled* with a number i from 1 to n and is *oriented* from a_i^{-1} to a_i^{+1} .

Let us denote by \mathcal{AMC} the set of those marked chord diagrams which are *analytic*, i.e. which arise from some analytic curve $F(x, y) = 0$ which does not contain the y -axis. Note that the analyticity of a marked diagram depends neither on the orientation of the chords nor on the labeling. The role of the labelings and orientations is simply to give the relevant information about which marked diagram is inserted in each chord, and in which way.

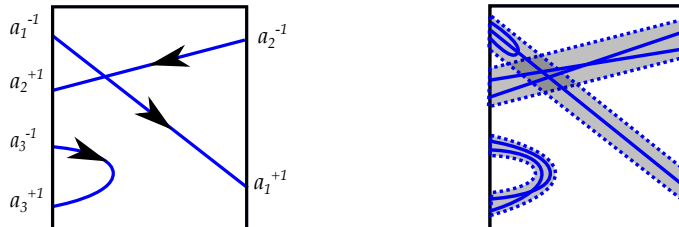


Figure 19: Inserting marked chord diagrams.

Let w be some analytic marked chord diagram with n chords. Given n analytic marked chord diagrams w_1, \dots, w_n , with k_1, \dots, k_n chords, define the action of w on (w_1, \dots, w_n) in the following way. Draw w and thicken each chord a_i of w , creating rectangles. Use the a_i^{-1} 's and the a_i^{+1} 's as the left and right sides of these new rectangles. Now, insert w_1, \dots, w_n in these rectangles respecting the labels and the orientations. Rename the chords, from 1 to $k_1 + k_2 + \dots + k_n$ using the lexicographic ordering. The result is another marked analytic chord diagram since this operation corresponds to the previously described insertion of analytic curves.

This defines an operad structure on \mathcal{AMC} though we will not make use of it in this paper.

1.8 Let us bound the number of chord diagrams

It would be great to have some precise information on the number a_n of analytic diagrams with n chords. For instance, an explicit formula for the generating series $\sum a_n t^n$ would give the exact exponential growth rate of a_n . Unfortunately, we were not able to compute this function. In this subsection, we show at least that the fundamental lemma provides a reasonable bound.

Consider a finite planar rooted *binary* tree. Equip each of its interior nodes (including the root) with one of the six examples of marked diagrams with two chords represented in Figure 20.

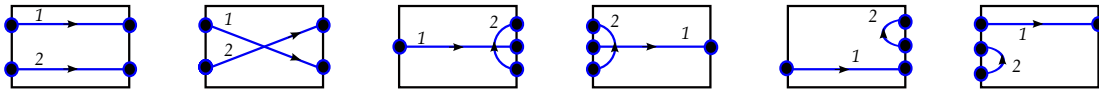


Figure 20: Six marked diagrams with two chords.

By recursive insertions of the diagrams of the siblings into the diagram of their parent, this produces a marked analytic diagram and hence an analytic diagram, forgetting the labels, the orientations, and the two sides of the square.

We claim that all analytic diagrams with $n \geq 2$ chords are produced by this recipe. This is true for $n = 2$ since both diagrams with 2 chords (linked or not linked) appear when one forgets the marking in the six examples. Now, let w be some analytic diagram with $n + 1$ chords, and apply the fundamental lemma. Therefore, we find $\dots aa \dots$, $\dots ab \dots ba \dots$, $\dots ab \dots ab \dots$ or $\dots b \dots aba \dots$ in the diagram. In the case of aa , call b the letter which comes *before* a in the cyclic order. Our algorithm deletes a and produces an analytic diagram \bar{w} with n chords, for which we can apply the induction. This means that \bar{w} can be decorated with labels and orientations in such a way that it is produced by a binary tree, as above. Our diagram w is obtained from \bar{w} by replacing one chord by two chords. It is easy to check that our six examples are sufficient to realize this duplication using an insertion in the operad. Hence w is constructed from a binary tree with $n + 1$ leaves with the same recipe.

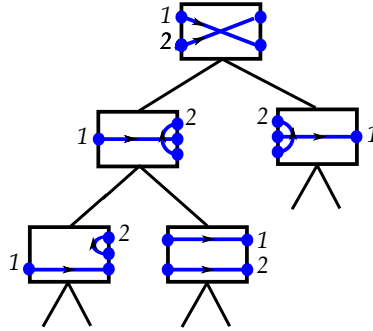


Figure 21: Recursive construction of marked diagrams.

A rooted binary tree with n leaves has $n - 1$ interior nodes (including the root) so that there are 6^{n-1} possible labels on the interior nodes. The number of planar binary trees with n leaves is given by the $(n - 1)$ -th Catalan number. Therefore, we get the following rough estimate.

Theorem. *The number a_n of analytic chord diagrams with n chords is less than 6^{n-1} times the $(n - 1)$ -st Catalan number C_{n-1} .*

The growth of Catalan numbers is well known: $\frac{1}{n} \log C_n$ converges to $\log 4$ as n goes to infinity. Therefore

$$\limsup \frac{1}{n} \log a_n \leq \log(24).$$

2 Analytic chord diagrams: interlace graphs

Let us begin by a definition.

Definition. The *interlace graph* of a chord diagram is the graph whose vertices are the chords and such that an edge connects two chords if they intersect.

Not every graph comes from a chord diagram and a graph might come from several chord diagrams. Nevertheless, the interlace graphs coming from analytic chord diagrams turn out to be easy to analyze. The icing on the cake is that these graphs have been introduced forty years ago in a totally different context and are very well understood. Thanks to this new perspective, we will get the complete description of basic non-analytic chord diagrams.

2.1 Polynomial interchanges: permutation graph

Before discussing analytic chord diagrams, we come back to the easy example of polynomial interchanges, in order to describe our general strategy.

Definition. The *permutation graph* $G(\pi)$ associated to a permutation π of $\{1, \dots, n\}$ has $\{1, \dots, n\}$ as vertices, and an edge connects i and j if π reverses the order of (i, j) .

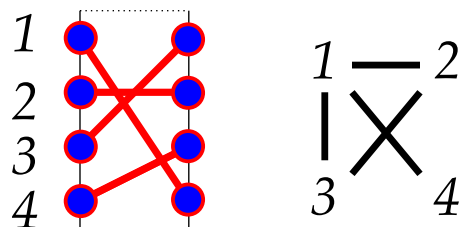


Figure 22: A permutation and its graph.

Note that a permutation of $\{1, \dots, n\}$ defines a (marked) chord diagram with n chords, as shown in Figure 23. The graph $G(\pi)$ is nothing more than the interlace graph of this chord diagram.

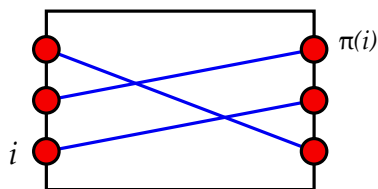


Figure 23: A permutation and its graph.

Recall that for any polynomial interchange, one can find two consecutive integers with consecutive images. The corresponding chords have therefore the property that any chord intersecting one of them intersects the other. In terms of the graph $G(\pi)$, this suggests the following definition.

Definition. Two vertices x, y in a graph are called *twins* if they have the same neighbors (different from x or y). They are called *true or false twins*² depending on the existence of an edge connecting them.

²We are not responsible for the terminology which is classical and more convenient than dizygotic (or fraternal) twins. It is also closer to the French “faux jumeaux”.



Figure 24: False and true twins.

Twins can be merged in a single vertex, producing a smaller graph with one vertex less. The graph $G(\pi)$ coming from a polynomial interchange π contains at least two twins, corresponding to two consecutive integers $i, i + 1$ such that $\pi(i + 1) = \pi(i) \pm 1$. Merging the twins in the graph amounts to merging the two elements $i, i + 1$. We know that polynomial interchanges are characterized by the fact the iteration of this merging procedure eventually leads to the trivial permutation with $n = 1$.

Definition. A finite graph is called a *cograph* if it can be reduced to a trivial 1-vertex graph by merging twins successively.

The terminology “cograph” comes from the fact that the *complement* of a cograph is also a cograph. A graph G and its complement \overline{G} have the same vertices and two vertices are adjacent in \overline{G} if and only if they are not adjacent in G .

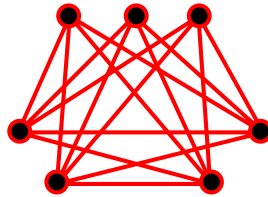


Figure 25: A cograph.

Proposition. *A permutation is a polynomial interchange if and only if its permutation graph is a cograph.*

We just explained why the permutation graph of a polynomial interchange is a cograph. To prove the converse, it suffices to show that if $G(\pi)$ is a cograph, there are two consecutive integers with consecutive images. The proof is by induction on n . If $i < j$ are false (resp. true) twins, the image by π of the interval $[[i, i + 1, \dots, j]]$ is $[[\pi(i), \pi(i + 1), \dots, \pi(j)]]$ (resp. $[[\pi(j), \pi(j + 1), \dots, \pi(i)]]$). If $j \geq i + 2$, the image $\pi([[i, \dots, j - 1]])$ is also an interval and we apply the induction hypothesis to the restriction of π to $[[i, \dots, j - 1]]$ so that one finds two consecutive integers with consecutive images. \square


Cographs have been introduced in the 1970’s under different names (D^* -graphs, hereditary Dacey graphs, and 2-parity graphs; see [2] for references). They are not very difficult

to describe. We list some of their properties whose (elementary) proofs are left to the reader.

In what follows, all graphs are finite, with no loops and no multiple edges. A connected graph defines a metric space on its set of vertices. The distance between two vertices is, by definition, the length of the shortest path connecting them.

A subgraph H of a graph G is *induced* if any edge of G connecting two vertices of H is also an edge of H .

Theorem. *The following properties of a finite graph G are equivalent.*

1. G is a cograph.
2. G is the permutation graph of some polynomial interchange.
3. Any connected induced subgraph of G has a diameter at most 2.
4. There is no induced subgraph isomorphic to  (denoted P_4).

Note that the permutation graphs associated to the forbidden Kontsevich permutations $(2, 4, 1, 3)$ and $(3, 1, 4, 2)$ are both isomorphic to P_4 . Conversely a permutation whose graph is isomorphic to P_4 is one of these two permutations. The previous theorem can therefore be translated as the following.

Theorem. *Let $n \geq 2$ be some integer. A permutation π of $\{1, 2, \dots, n\}$ is a polynomial interchange if and only if it does not “contain” one of the two forbidden permutations, i.e. if there do not exist four indices $1 \leq i_1 < i_2 < i_3 < i_4 \leq n$ such that $\pi(i_2) < \pi(i_4) < \pi(i_1) < \pi(i_3)$ or $\pi(i_3) < \pi(i_1) < \pi(i_4) < \pi(i_2)$.*

2.2 Collapsible graphs

Starting from a tree, one can strip off its leaves one by one until it has been stripped completely naked. Let us say that a vertex in a graph is *pendant* if it is adjacent to a unique vertex. Any tree can be constructed by successive additions of pendant vertices, starting with the tree with only one vertex.

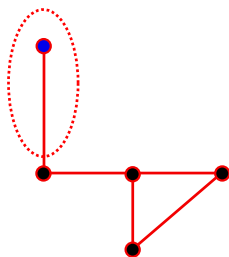


Figure 26: Pendant vertex.

Definition. A finite graph is *collapsible* if it can be reduced to a 1-vertex graph by applying two kinds of elementary operations: deleting a pendant vertex and merging twins.

One can express the same thing in the opposite way. Start with the trivial graph with one vertex and apply two kinds of operations: adding a pendant vertex or creating a pair of twins. The second operation simply consists in duplicating a vertex and connecting the newly born twin to the rest of the graph as the original vertex was. Then, decide if you want true or false twins.

The key point is the following.

Proposition. A chord diagram is analytic if and only if its interlace graph is collapsible.

This will follow from the algorithmic description of analytic diagrams given in the previous section.

Before the proof, let us make an elementary remark, as an appetizer.

Let w be a diagram and A be a subset of its $2n$ letters on the circle. We will say that A is *stable* under w if any chord with one end in A has its other end in A . Said differently A is a sub-chord diagram w_A of w .

Suppose that there is an *interval* A which is stable under w , and let B be its complement. Clearly the interlace graph $G(w)$ of w is the disjoint union of the graphs $G(w_A)$ and $G(w_B)$ of w_A and w_B . It follows that $G(w)$ is collapsible if and only if $G(w_A)$ and $G(w_B)$ are collapsible. Our algorithm shows that if w_A, w_B are analytic so is w . Conversely, if w is analytic, so are their sub-diagrams w_A and w_B .

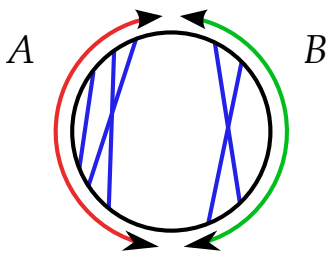
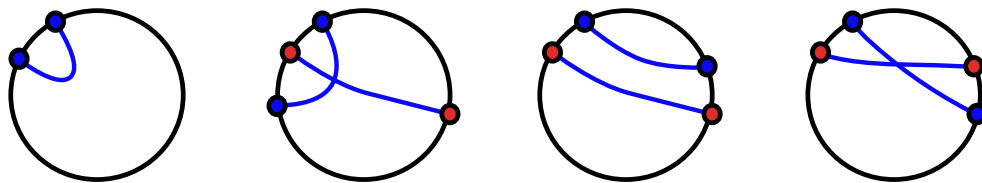


Figure 27: $G(w) = G(w_A) \sqcup G(w_B)$.

Let us prove now the proposition.

Start with an analytic diagram w . If two chords of w are parallel or antiparallel, the associated vertices in the interlace graph are twins and our algorithm merges them. A pitchfork gives a pendant vertex in the graph and the algorithm deletes the short chord and keeps the handle. A solitary chord defines an isolated vertex, which is removed by the algorithm. It follows that the interlace graph associated to an analytic diagram is collapsible.



For the converse, we show that every diagram w whose interlace graph $G(w)$ is collapsible contains a solitary chord, or a pitchfork, or a pair of parallel or antiparallel chords.

A solitary chord in a diagram corresponds to two consecutive identical letters $\dots aa \dots$ in the cyclic word. A pitchfork corresponds to a subword of the form $\dots aba \dots$. A pair of parallel (resp. antiparallel) chords corresponds to $\dots ab \dots ba \dots$ (resp. $\dots ab \dots ab \dots$).

Our proof will be by contradiction. We use the symbol $\not\equiv$ to mean “contradiction”. Consider a possible counterexample w to the previous assertion *with a minimal number of chords*. So $G(w)$ is collapsible and w contains no solitary chord, no pair of parallel or antiparallel chords, and no pitchfork.

Since $G(w)$ is collapsible, there is a vertex α which is either isolated, or pendant, or is part of a pair of twins. Let \bar{w} be the diagram obtained by deleting α from w . Of course $G(\bar{w})$ is collapsible so that, by minimality, \bar{w} contains a subword $\dots aa \dots$ or $\dots aba \dots$, or $\dots ab \dots ba \dots$, or $\dots ab \dots ab \dots$. The problem is that these are subwords of \bar{w} and not of w , which also contains two copies of the letter α , which could sneak into the above subwords. Note that by minimality any interval which is stable under w is either empty or everything.

A priori:

- 0, 1 or 2 letters α could sneak in the subword, i.e. in non dotted intervals
- the subword of \bar{w} could correspond to a solitary chord, or a pitchfork, or to a pair of parallel or antiparallel chords,
- α could be isolated, pendant or twin, true or false, in $G(w)$.

That makes $3 \times 4 \times 4$ cases to examine! Fortunately, many cases can be studied simultaneously.

1/ If no letter α sneaks into the above subwords, there is no problem: our solitary chord or pitchfork, or pair of parallel or antiparallel chords in \bar{w} have the same property for w $\not\equiv$.

2/ If α is isolated in $G(w)$, this means that no chord intersects α . Therefore α decomposes the circle in two stable intervals, of which one has to be empty $\not\equiv$.

3/ If two letters α sneak in, they cannot occur as consecutive letters since that would force the chord α to be solitary in w $\not\equiv$. Still in this case, we have to look at

$$\dots a\alpha b\alpha a \dots, \text{ or } \dots a\alpha b \dots b\alpha a \dots, \text{ or } \dots a\alpha b \dots a\alpha b \dots$$

This produces respectively a pitchfork (α, b) , a pair of parallel chords (α, b) or antiparallel chords (α, b) in w $\not\equiv$.

4/ Inserting one α in a solitary chord yields $\cdots a\alpha a \cdots$ which produces a pitchfork in w with handle $\alpha \notin$.

So far, we did not use the fact that α is pendant or twin. This will be used in the remaining cases, when a single α enters in a pitchfork or a pair of parallel or antiparallel chords of \bar{w} .

If α is pendant, let β be the only chord in w intersecting α . If α has twin siblings, we denote β one of them. The two chords α, β determine four intervals in the circle, *excluding* α, β , that will be called *sectors*. If α is pendant, the union of two sectors which are on the same side of α is stable. If α, β are twins, the union of opposite sectors is stable.

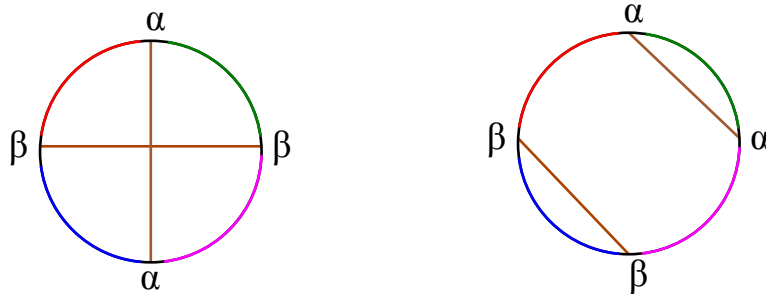


Figure 28: Linked or unlinked.

5/ Suppose now that a single letter α enters $\cdots aba \cdots$, or $\cdots ab \cdots ba \cdots$, or $\cdots ab \cdots ab \cdots$ and that one of the two letters a, b is equal to β . Then, the two letters α and β are consecutive in w . This implies that one of the sectors is empty.

5-1 In the pendant case, this implies that the other sector, on the same side of α , is a stable interval and therefore also empty. So, α, β is a pitchfork in w with handle $\beta \notin$.

5-2 In the twin case, this implies that the opposite sector is a stable interval, and therefore empty. So α, β is a pair of parallel or antiparallel chords in $w \notin$.

6/ Finally, suppose that a single letter α enters $\cdots aba \cdots$, or $\cdots ab \cdots ba \cdots$, or $\cdots ab \cdots ab \cdots$ and that none of the two letters a, b is equal to β .

6-1 Assume that (a, b) are parallel or antiparallel chords in \bar{w} . Inserting one α in ab yields $\cdots a\alpha b \cdots ba \cdots$ or $\cdots a\alpha b \cdots ab \cdots$. Since the letters ba (and ab) are consecutive, these two occurrences are on the same side of the chord α . And since the letters $a\alpha b$ are consecutive in w , it follows that these other two occurrences of a and b are on different sides of α . Hence the chord α intersects only one of the chords a, b .

6-1-1 If α is pendant, this forces a or b to be equal to $\beta \notin$.

6-1-2 If α, β are twins, this is not possible \notin .

6-2 Assume that (a, b) is a pitchfork in \bar{w} . Inserting a letter α in $\cdots aba \cdots$ yields $\cdots a\alpha ba$ or $\cdots ab\alpha a$ so that the chord α should intersect the chord a .

6-2-1 If α is pendant, this forces $a = \beta \notin$.

6-2-2 If (α, β) are twins, then β (which is not a) also intersects a . Since the letters $a\alpha b a$ are consecutive, we must have $\beta = b \notin$.

This finishes the proof.

Ouf ! \square

Now, we have to understand the nature of collapsible graphs.

2.3 Distance hereditary and treelike graphs

Collapsible graphs have been defined by several authors forty years ago, under different names, with very different motivations. We will see that these graphs are very close to being trees.

Howorka [11] defined *distance hereditary graphs* in 1977.

Definition. A finite graph G is *distance hereditary* if for every connected induced subgraph $H \subset G$, the distance between two vertices of H in H is equal to the distance between the same vertices in G .

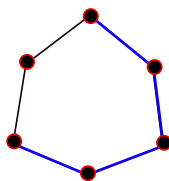


Figure 29: A cycle of length ≥ 5 is not distance hereditary.

For instance, a tree is *distance hereditary* and a cycle of length at least 5 is not. It suffices to choose H as the induced subgraph defined by a path inside the cycle whose length is greater than one half of the length of the cycle.

Let us introduce *metric graphs* along with a property characterizing the subspaces of *metric trees*. A weakened version of this property will give an alternative description of our distance hereditary graphs. Consider a finite graph and choose some length for each edge, which could be any positive real number. Define the length of a path as the sum of the lengths of its edges and the distance between two vertices as the smallest length of a path connecting them. One speaks of a *metric graph*.

We are looking for a characterization of metric spaces (usually called *metric trees*) arising in this way from *trees*. Here is the answer. Let (V, d) be a finite metric space. Choose four points x_1, x_2, x_3, x_4 in V and compute the sums of the lengths of the three pairs of diagonals:

$$d(x_1, x_2) + d(x_3, x_4); d(x_1, x_3) + d(x_2, x_4); d(x_1, x_4) + d(x_2, x_3).$$

Let s (resp. m , l) be the smallest (resp. medium, largest) of these three numbers: $s \leq m \leq l$. It turns out that a finite metric space is isometric to a subset of a metric tree if and only if $m = l$ for every quadruple of points. This is not difficult to prove but the lazy reader might find the proof in this short paper [3].

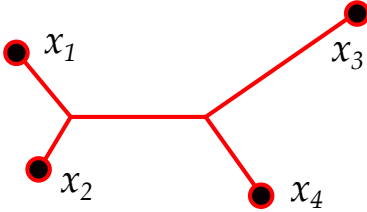


Figure 30: A quadrangle in a tree

We should be careful. A graph, where all edges have length 1, can be isometric to a subset of a metric tree without being itself a tree. Look at the example in Figure 31. In graph theory, those graphs are called *block graphs*.

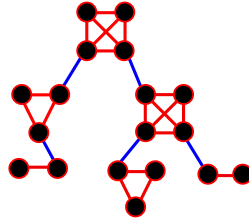


Figure 31: A block graph.

In order to construct them, start with a tree, delete some of its vertices and replace them by *cliques*, i.e. finite graphs where all pairs of vertices are adjacent, as in the Figure. This is indeed a characterization of block graphs (see [10]).

A metric space (E, d) is *geodesic* if for every pair of points (x, y) there exists an *isometric* embedding $i : [0, d(x, y)] \rightarrow E$ such that $i(0) = x$ and $i(d(x, y)) = y$.

In the 1980's, Gromov developed a geometric theory for *hyperbolic spaces* which had a very strong influence on combinatorial and geometric group theory. The definition is the following. A metric space (E, d) is called *hyperbolic* if there exists some $\delta \geq 0$ such that for every quadruple of points as above, m and l are “almost equal”, i.e. $l - m \leq \delta$. Note that any finite metric space is trivially hyperbolic (for δ sufficiently big) so that this concept is only relevant for geometry *in the large*. There are many equivalent formulations of this property, the most popular (for geodesic metric spaces) being that all geodesic triangles are *slim*. Consider three points x, y, z and choose three geodesics $[x, y], [x, z], [y, z]$ connecting them. Every point in $[x, y]$ should be at some uniformly bounded distance from the union $[x, z] \cup [y, z]$, independently of the choice of x, y, z (Figure 32).

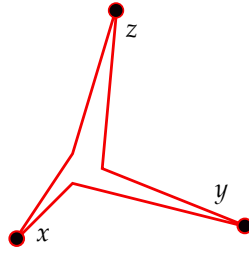


Figure 32: A slim triangle

This concept is remarkably robust. For instance, the universal cover of a negatively curved compact Riemannian manifold is hyperbolic. These metric spaces are well approximated by trees, in a quantitative way (see for instance [7]).

In 1986, Bandelt and Mulder published a paper [1] proposing purely metrical characterizations of distance hereditary graphs, close to Gromov's hyperbolicity conditions.

Definition. A finite graph G is *treelike* if for every 4-tuple of vertices x_1, x_2, x_3, x_4 two of the following three numbers are equal:

$$d(x_1, x_2) + d(x_3, x_4); \quad d(x_1, x_3) + d(x_2, x_4); \quad d(x_1, x_4) + d(x_2, x_3).$$

For instance, one can show easily that a treelike graph is hyperbolic in the sense of Gromov with $\delta=2$.

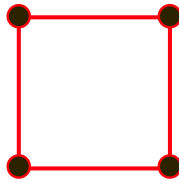


Figure 33: A cycle of length 4 is treelike but not a tree.

All these definitions turn out to be equivalent, and correspond to

Theorem. *Let G be a finite graph. The following properties are equivalent.*

1. G is the interlace graph of some analytic chord diagram,
2. G is collapsible,
3. G is distance hereditary,
4. G is treelike,
5. G does not contain a cycle of length at least five, or a house, a gem, or a domino (see Figure 34), as an induced subgraph.

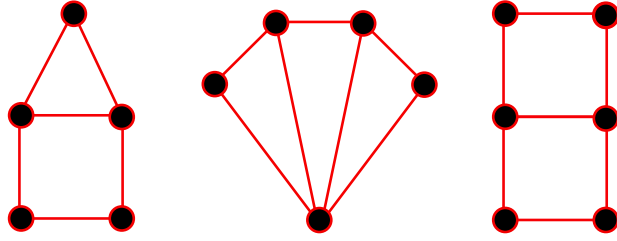
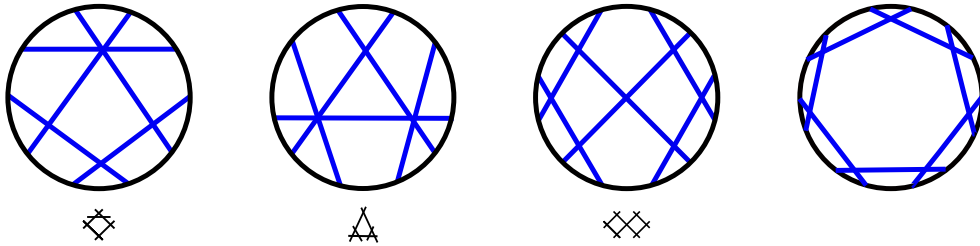


Figure 34: The *house*, the *gem* and the *domino*

All the equivalences in the previous theorem (except of course the first item) are proved in the above mentioned papers. However, we will soon propose some elementary proofs.

It is now time to harvest the fruits of our labor and to get a very simple description of analytic chord diagrams.



It should not be surprising that the interlace graphs of \diamond , \triangle , $\diamond\diamond$ are the house, the gem, and the domino. In fact, one can easily show that $\diamond\diamond$, \diamond , \triangle are the only chord diagrams whose interlace graphs are the house, the domino and the gem.

In the same way, we have already described the non-analytic chord diagram C_n defined by $\mathbb{Z}/2n\mathbb{Z}$ ($n \geq 5$) where there is a chord connecting $2k$ and $2k + 3$ (for $k = 1, \dots, n$). Its interlace graph is a cycle of length n . It also clear that C_n is the only chord diagram whose interlace graph is a cycle of length n .

Finally, note that a sub-chord diagram defines an induced subgraph in the interlace graph. Therefore, we get a very satisfactory description of analytic chord diagrams, which is the main result of this paper.

Theorem. *A chord diagram is analytic if and only if it does not contain $\diamond\diamond$, \diamond , \triangle or C_n ($n \geq 5$) as a sub-chord diagram.*

Note the complete analogy with our characterization of *polynomial interchanges* as those permutations which do not contain Kontsevich's examples $(2, 4, 1, 3)$ and $(3, 1, 4, 2)$.

2.4 Some proofs

We present now the proofs of the equivalences of the definitions in the previous subsection. They are mostly elementary. In this specific case, it was probably more challenging to find the significant definitions than to prove their equivalence.

No $\triangle \diamond \square \circlearrowleft$ as an induced subgraph \implies Distance hereditary.

Let H be an induced subgraph of a graph G . Connect two vertices p, q of H at distance n in H by a path $c = (x_0, x_1, \dots, x_n)$ (with $p = x_0$ and $q = x_n$) in H . Two vertices x_i, x_j are adjacent if and only if i, j are consecutive since otherwise there would be a shortcut. In other words, the path c is induced in G . It follows that in order to show that a graph is distance hereditary, we should prove that the distance between the endpoints of any induced path in G is equal to the length of the path.

Suppose that no $\triangle \diamond \square \circlearrowleft$ is induced in G . Choose an induced path $c_1 = (x_0, x_1, \dots, x_n)$ and let us show, by induction on n , that the distance in G between x_0 and x_n is exactly n .

Connect x_0 to x_n by a shortest path $c_2 = (y_0, y_1, \dots, y_l)$ in G (with $y_0 = x_0$ and $y_l = x_n$). Of course, c_2 is also induced and $d(y_0, y_i) = i$ for $0 \leq i \leq l$. By the induction hypothesis, $d(x_0, x_i) = i$ for $0 \leq i \leq n-1$ so l is equal to $n-2, n-1$ or n . We must show that the first two cases are not possible. Suppose that $l = n-2$ or $n-1$. By induction hypothesis, we can assume that the two paths c_1, c_2 only intersect at their endpoints: any other intersection point could be used as the starting point of shorter paths c'_1 and c'_2 .

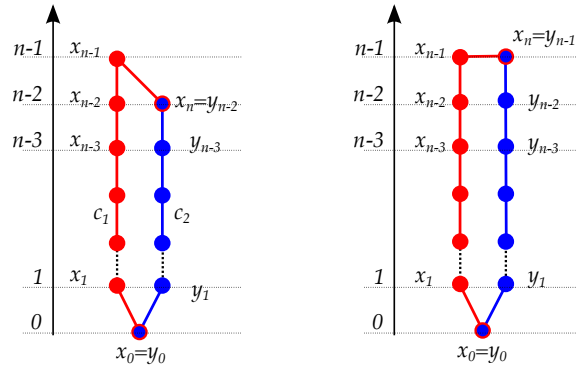


Figure 35: Checking heredity...

Draw a picture in the plane in such a way that the height of a vertex of c_1 or c_2 is the distance from x_0 . The cases $l = n-2$ and $l = n-1$ are pictured in Figure 35. Vertices of c_1 are red and vertices of c_2 are blue. The union of c_1 and c_2 defines a cycle c in G . The length of c is at least 5. The cycle c cannot be induced since there is no induced cycle of length ≥ 5 . Therefore there must exist diagonals connecting vertices of c_1 with vertices of c_2 .

By the triangle inequality, the height difference of the two endpoints of a diagonal can only be $-1, 0, 1$. Moreover, diagonals connect points of different colors. Let us order the diagonals (x_i, y_j) from top to bottom, i.e. (x_i, y_j) is before $(x_{i'}, y_{j'})$ if $j > j'$ or $j = j'$ and $i > i'$.

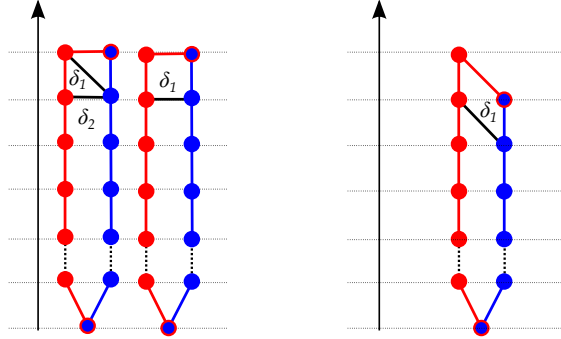


Figure 36: No $\triangle \diamond \square \boxtimes \bowtie \implies$ distance hereditary

Now, try to construct the ladder, one diagonal at a time. The rule of the game is the following. You have to draw an ordered sequence of diagonals $\delta_1, \delta_2, \dots$ respecting the conditions above and without creating any induced $\triangle \diamond \square \boxtimes \bowtie$. Note that the diagonal δ_k together with the part of c which is above it defines a cycle. Any chord in this cycle has to be one of the previously chosen chords $\delta_1, \dots, \delta_{k-1}$.

In the case $l = n - 1$, there are only two possibilities for the first diagonal δ_1 . It could be (x_{n-1}, y_{n-2}) or (x_{n-2}, y_{n-2}) . In the case $l = n - 2$ there is only one possibility for δ_1 . Then, try to select the second diagonal δ_2 , avoiding $\triangle \diamond \square \boxtimes \bowtie$. Only one of the three choices of δ_1 allows you to do so. Finally, try to draw the third diagonal, in the only case where you could draw δ_1, δ_2 . It is not possible to continue without creating one of the forbidden graphs. \square

Distance hereditary \implies Collapsible.

We only need to show that G has a pendant vertex or a pair of twins. Indeed deleting the former or merging the latter preserves the distance hereditary property so that we can repeatedly collapse the graph to a single vertex.

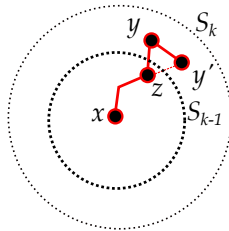


Figure 37: hereditary \implies Collapsible

Choose some vertex x in a connected distance hereditary graph G and look at the largest k such that the sphere S_k in G of radius k and centered on x is non-empty. Let

C be a connected component of S_k . If C contains only one element, then it is a pendant vertex in G .

Otherwise, choose two vertices y, y' in C which are adjacent in C . Choose a vertex z adjacent to y and at distance $k - 1$ from x . Choose a chain c of length $k - 1$ from x to z and call c' the chain of length $k + 1$ obtained by adding the edge between z and y and from y to y' . Since the distance between x and y' is exactly k , this chain cannot be induced and y' has to be adjacent to z . This implies that two points in C are simultaneously adjacent or not to any point z at distance $k - 1$ from x . It follows that C cannot contain an induced path P_4 of length 3, since together with z , it would produce a gem in G , which is not distance hereditary. Hence C is a cograph and in particular contains a pair of twins. By the above observation, two twins in C are twins in G . \square

Collapsible \implies *Treelike*.

Easy by induction. Take four points in a graph, delete a pendant vertex or merge two twins. One of the four points might be the vertex which has been removed. If this is the case replace it by the other end of the removed edge. Look at the corresponding points in the stripped graph (taking into account for instance the fact that two of our four points could be the two twins which have been merged). Apply the induction hypothesis. \square

Treelike \implies No induced $\triangle \diamond \boxplus \boxminus$.

Obvious since one checks easily that none of these examples of graphs are treelike. \square

2.5 Appendix : completely decomposable graphs

This subsection is some kind of bonus but it can shed some light on the structure of collapsible graphs.

Graphs that can be stripped to a point by deleting only pendant vertices are trees. Graphs that can be stripped to a point by merging only pairs of twins are cographs. Collapsible graphs should not be far from being trees. This is indeed true as explained now.

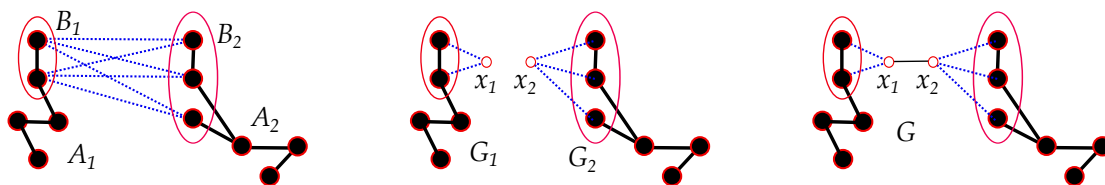


Figure 38: Splitting a graph

Let G be a finite connected graph. Suppose that its vertices have been partitioned in two parts A_1 and A_2 . Let $B_1 \subset A_1$ (resp. $B_2 \subset A_2$) the set of vertices of A_1 (resp. A_2) which are adjacent to some vertex in A_2 (resp. A_1). Suppose that every element of

B_1 is adjacent to every element of B_2 and that those are the only edges between A_1 and A_2 . This condition is trivially satisfied if A_1 or A_2 contains less than one element, so we assume that A_1 and A_2 contain at least two elements each. In this situation, the graph G is called *decomposable* and the partition A_1, A_2 is a *split*. In order to keep track of this decomposition, let us create two graphs G_1, G_2 in the following way. The set of vertices of G_1 (resp. G_2) is A_1 plus one extra vertex x_1 (resp. x_2) called the *control vertex*. As for the edges of G_1 (resp. G_2), choose the edges of G plus extra edges connecting x_1 (resp. x_2) to all elements of B_1 (resp. B_2).

The graph G can be reconstructed from (G_1, x_1) and (G_2, x_2) by an elementary *join* construction. Notice that the control points x_1, x_2 are not vertices of G : they are only useful to define the edges connecting the two parts.

Note that when A_2 contains two elements, as in Figure 39, the graph G has a pendant vertex or a pair of twins.

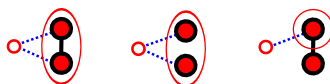


Figure 39: A_2 contains 2 elements.

Hammer and Maffray [9] introduced in 1987 *completely decomposable* graphs, analog to trees. A finite connected graph is a tree if and only if every induced connected subgraph contains a *cut edge*, i.e. an edge that disconnects it.

Definition. A finite graph is *completely decomposable* if every induced connected subgraph with at least four vertices is decomposable.

It is not hard to prove that completely decomposable (connected) graphs are precisely the collapsible (connected) graphs. Indeed in the join construction, if G_1 and G_2 are collapsible, the same is true for G , so that completely decomposable graphs are collapsible, by induction.

Conversely, we have seen that a pendant vertex or a pair of twins gives rise to a decomposition. Therefore collapsible graphs are decomposable and even completely decomposable, since an induced subgraph of a collapsible graph is collapsible (with the distance hereditary definition).

In order to give a precise description of completely decomposable graphs, let us state first the important *split decomposition* theorem for *general* connected finite graphs.

If a finite connected graph G is decomposable, consider it as the join of G_1 and G_2 as before. Then, try to decompose G_1 and G_2 , and so on, until the resulting graphs become non-decomposable. The final result of this decomposition into “elementary pieces” can be conveniently described by a *graph-labeled tree*, as explained below. It consists of a tree T where each internal node x is equipped with a connected finite graph G_x which is either

a clique, a star or indecomposable (no splits). Moreover some bijection has been chosen between the vertices of G_x and the edges getting out from the node x in T . Assume that the valency of each node is at least 3. Given such a structure, we construct a graph $G(T)$ which is a “composition of the G_x ’s controlled by T ”. The definition is the following.

The vertices of $G(T)$ are the leaves of T . In order to understand the edges of $G(T)$, let us just draw a picture, inspired by the paper of Gioan and Paul [8] who introduced this concept of *graph-labeled tree*. We see a tree with 16 leaves and 6 internal nodes, in pink. The associated graph, with 16 vertices, is drawn on the right.

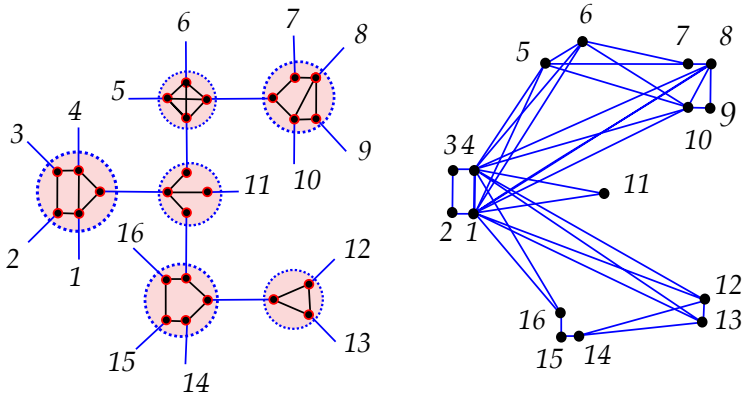


Figure 40: Graph-labeled tree and its composition

Choose two leaves of T and connect them by the shortest path in the tree. For each node x which is visited by this path, there is an entrance edge and an exit edge. In turn, these two edges define two vertices of G_x . Two vertices of $G(T)$, that is two leaves of T , are adjacent in $G(T)$ if, for every node x visited by this path, the two corresponding vertices of G_x are adjacent in G_x . The vertices of the G_x ’s generalize the two control vertices, as in the simple case where T contains only one edge.

The main result, proved by Cunningham and Edmonds [4] in 1980 (and reformulated by Giona and Paul), is that *any finite connected graph is obtained by such a construction in which the G_x ’s are elementary, in an essentially unique way*. The *existence* of this splitting is easy. The hard part is the “essential uniqueness” that we don’t define since we will not need it.

Let us come back to completely decomposable graphs. In this special case the G_x ’s must *have at most 3 vertices*. Indeed, they are indecomposable and induced subgraphs of G , so that the claim follows from the definition of complete decomposability. This gives a fairly precise geometric description of completely decomposable graphs. Take a tree such that every node has valency 3. For each node, choose a connected graph with 3 vertices (not that many choices!), and construct a graph-labeled tree as in Figure 41. All completely decomposable graphs are produced in this way.

This is absolutely no surprise. Indeed look at the third picture, showing three small graphs with three leaves. When you hook one of them to some (blue) leaf of a graph-labeled

tree T , you get another graph-labeled tree T' with one more leaf. If you examine the effect on the associated graph $G(T)$, you see that you have split a vertex in a pair of twins (true or false) or you have created a pendant vertex, depending on the cases. We are back to the original definition of collapsible graphs as graphs that can be constructed from a point by successive introductions of twins or of pendant vertices.

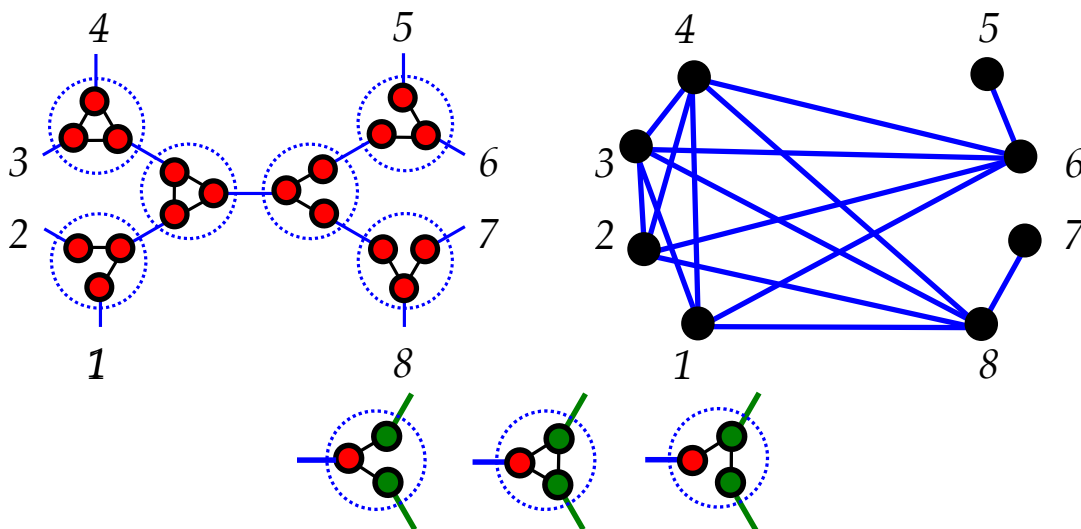


Figure 41: A collapsible graph

References

- [1] H.-J. Bandelt and H. M. Mulder. Distance-hereditary graphs. *J. Combin. Theory Ser. B*, 41(2):182–208, 1986.
- [2] A. Brandstädt, V. B. Le, and J. P. Spinrad. *Graph classes: a survey*. SIAM Monographs on Discrete Mathematics and Applications. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1999.
- [3] P. Buneman. A note on the metric properties of trees. *J. Combinatorial Theory Ser. B*, 17:48–50, 1974.
- [4] W. H. Cunningham and J. Edmonds. A combinatorial decomposition theory. *Canad. J. Math.*, 32(3):734–765, 1980.
- [5] É. Ghys. Intersecting curves (variation on an observation of Maxim Kontsevich). *Amer. Math. Monthly*, 120(3):232–242, 2013.
- [6] É. Ghys. *A singular mathematical promenade*. ENS Editions, 2017.

- [7] É. Ghys and P. de la Harpe, editors. *Sur les groupes hyperboliques d'après Mikhael Gromov*, volume 83 of *Progress in Mathematics*. Birkhäuser Boston, Inc., Boston, MA, 1990.
- [8] E. Gioan and C. Paul. Split decomposition and graph-labelled trees: characterizations and fully dynamic algorithms for totally decomposable graphs. *Discrete Appl. Math.*, 160(6):708–733, 2012.
- [9] P. L. Hammer and F. Maffray. Completely separable graphs. *Discrete Appl. Math.*, 27(1-2):85–99, 1990. Computational algorithms, operations research and computer science (Burnaby, BC, 1987).
- [10] F. Harary. A characterization of block-graphs. *Canad. Math. Bull.*, 6:1–6, 1963.
- [11] E. Howorka. A characterization of distance-hereditary graphs. *Quart. J. Math. Oxford Ser. (2)*, 28(112):417–420, 1977.
- [12] A. Stoimenow. On the number of chord diagrams. *Discrete Math.*, 218(1-3):209–233, 2000.

ÉTIENNE GHYS AND CHRISTOPHER-LLOYD SIMON

Unité de Mathématiques Pures et Appliquées
de l'École normale supérieure de Lyon
U.M.R. 5669 du CNRS
46, Allée d'Italie
69364 Lyon Cedex 07- France
etienne.ghys@ens-lyon.fr
christopher-lloyd.simon@ens-lyon.fr

April 15, 2018
