A monodromy criterion for existence of Neron models and a result on semi-factoriality

Orecchia, G.

Author: Orecchia, G.
Title: A monodromy criterion for existence of Neron models and a result on semi-factoriality
Issue Date: 2018-02-27

We find that the term $H^1(\mathcal{Z}_1, \mathcal{O}_{\mathcal{Z}_1}(n))$ vanishes using Mayer-Vietoris exact sequence and the fact that $H^1(\mathbb{P}^1_k, \mathcal{O}_{\mathbb{P}^1_k}(n)) = 0$. It follows that the restriction map $\mathrm{Pic}(\mathcal{Z}_{n+1}) \to \mathrm{Pic}(\mathcal{Z}_n)$ is an isomorphism. Since the sheaf $\mathcal{L}_{|\mathcal{Z}_{n+1}}$ restricts to the trivial sheaf on $\mathcal{Z}_n$, it is itself trivial, establishing the claim.

We obtain

$$\lim_n H^0(\mathcal{Z}_n, \mathcal{L}_{|\mathcal{Z}_n}) \cong \lim_n H^0(\mathcal{Z}_n, \mathcal{O}_{\mathcal{Z}_n}) \cong \lim_n (\pi_*\mathcal{O}_{\mathcal{Y}}) \otimes_{\mathcal{O}_{\mathcal{X}}} \mathcal{O}_p/\mathfrak{m}_p^n \cong \widehat{\mathcal{O}}_p$$

the second isomorphism coming again from the formal function theorem applied to $\mathcal{O}_{\mathcal{Y}}$ and the third coming from lemma 10.1. Finally, we obtain by composition with Φ an isomorphism

$$\lim_n (\pi_*\mathcal{L}) \otimes_{\mathcal{O}} \mathcal{O}_p/\mathfrak{m}_p^n \to \widehat{\mathcal{O}}_p$$

which induces an isomorphism $\pi_*\mathcal{L} \otimes_{\mathcal{O}} \mathcal{O}_p/\mathfrak{m}_p \to \mathcal{O}_p/\mathfrak{m}_p = k(p)$, as desired.

Now we drop the assumption of strict henselianity on the base, so let S be the spectrum of a discrete valuation ring. Let S' be the étale local ring of S with respect to some separable closure of the residue field of S. The cartesian diagram

$$\begin{array}{ccc} \mathcal{Y}_{S'} & \xrightarrow{f} & \mathcal{Y} \\ \downarrow{\scriptstyle \pi'} & & \downarrow{\scriptstyle \pi} \\ \mathcal{X}_{S'} & \xrightarrow{g} & \mathcal{X} \end{array}$$

has faithfully flat horizontal arrows, and $\mathcal{Y}_{S'} \to \mathcal{X}_{S'}$ is the blowing-up at $g^{-1}(p)$. Let $\mathcal{L}$ be a line bundle on $\mathcal{Y}$ as in the hypotheses. The restrictions of $f^*\mathcal{L}$ to the irreducible components of the exceptional fibre of π' have degree zero, hence $\pi'_*f^*\mathcal{L}$ is a line bundle. Moreover the canonical map

$$g^*\pi_*\mathcal{L} \to \pi'_*f^*\mathcal{L}$$

is an isomorphism, because g is flat. Hence $g^*\pi_*\mathcal{L}$ is a line bundle, and so is $\pi_*\mathcal{L}$ by faithful flatness of g. □

11 Graph theory

In this section we develop some graph-theoretic results that, together with the results of sections 9 and 10, will be needed to prove theorem 12.3.

11.1 Labelled graphs

Let $G = (V, E)$ be a connected, finite graph. For the whole of this section, we will just write "graph" to mean finite, connected graph. A *circuit* in G is a closed walk in G all of whose edges and vertices are distinct except for the first and last vertex. A *path* is an open walk all of whose edges and vertices are distinct.

A *tree* of G is a connected subgraph $T \subset G$ containing no circuit. A *spanning tree* of G is a tree of G containing all of the vertices of G, that is, a maximal tree of G. Given a spanning tree $T \subset G$, we call *links* the edges not belonging to T.

Let $n = |E|$, $m = |V|$. Given a spanning tree T, the number of links of T is easily seen to be $n - m + 1$. The number

$$r := n - m + 1$$

is called *nullity* of G and is equal to the first Betti number $\operatorname{rk} H^1(G, \mathbb{Z})$.

Fix a spanning tree $T \subset G$. For each link $c_1, \ldots, c_r$ of T, the subgraph $T \cup c_i$ contains exactly one circuit $C_i \subset G$. We call $C_1, \ldots, C_r$ *fundamental circuits* of G (with respect to T).

Let $(G, l) = (V, E, l)$ be the datum of a graph and of a labelling of the edges $l \colon E \to \mathbb{Z}_{\geq 1}$ by positive integers. We say that (G, l) is a $\mathbb{N}$-*labelled graph.*

11.2 Circuit matrices

Given a graph G, let $e_1, e_2, \ldots, e_n$ be its edges and $\gamma_1, \ldots, \gamma_s$ its circuits. Fix an arbitrary orientation of the edges of G, and an orientation of each circuit (that is, one of the two travelling directions on the closed walk).

Definition 11.1. The *circuit matrix* of G is the $s \times n$ matrix M_G whose entries a_{ij} are defined as follows:

$$a_{ij} = \begin{cases} 0 & \text{if the edge } e_j \text{ is not in } \gamma_i; \\ 1 & \text{if the edge } e_j \text{ is in } \gamma_i \text{ and its orientation agrees} \\ & \text{with the orientation of } \gamma_i; \\ -1 & \text{if the edge } e_j \text{ is in } \gamma_i \text{ and its orientation does not agree} \\ & \text{with the orientation of } \gamma_i. \end{cases}$$

Hence every row of M_G corresponds to a circuit of G and each column to an edge.

Now fix a spanning tree of G. Let $c_1, \dots, c_r$ be the corresponding links, where r is the nullity of G, and $C_1, \dots, C_r$ the associated fundamental circuits. Consider the $r \times n$ submatrix N_G of M_G given by singling out the rows corresponding to fundamental circuits. One can reorder edges and circuits so that the j-th column corresponds to the link c_j for $1 \leq j \leq r$ and that the i-th row corresponds to the circuit C_i. If we also choose the orientation of every fundamental circuit C_i so that it agrees with the orientation of the link c_i, the matrix N_G has the form

$$N_G = [\mathbb{I}_r | N']$$

where $\mathbb{I}_r$ is the identity $r \times r$-matrix and N' is an integer matrix.

Definition 11.2. The matrix N_G constructed above is called the *fundamental circuit matrix* of G (with respect to the spanning tree T).

It is clear that N_G has rank r.

Theorem 11.3 ([TS92], Theorem 6.7.)**.** *The rank of M_G is equal to the rank of N_G.*

Let now (G, l) be an $\mathbb{N}$-labelled graph. We generalize the definitions above to this case.

Definition 11.4. The *labelled circuit matrix* of (G, l) is the $s \times n$ matrix $M_{(G,l)}$ whose entries b_{ij} are defined as follows:

$$b_{ij} = \begin{cases} 0 & \text{if the edge } e_j \text{ is not in } \gamma_i; \\ l(e_j) & \text{if the edge } e_j \text{ is in } \gamma_i \text{ and its orientation agrees} \\ & \text{with the orientation of } \gamma_i; \\ -l(e_j) & \text{if the edge } e_j \text{ is in } \gamma_i \text{ and its orientation does not agree} \\ & \text{with the orientation of } \gamma_i. \end{cases}$$

The *labelled fundamental circuit (lfc) matrix* of (G, l) is the $r \times n$ matrix $N_{(G,l)}$ constructed from $M_{(G,l)}$ by taking only the rows corresponding to fundamental circuits with respect to a given spanning tree T.

We immediately see that

$$M_{(G,l)} = M_G \cdot L \text{ and } N_{(G,l)} = N_G \cdot L$$

where L is the diagonal square matrix of order n whose (i, i)-th entry is $l(e_i)$.

Example 11.5. Consider the $\mathbb{N}$-labelled graph (G, l) with oriented edges in fig. 1.

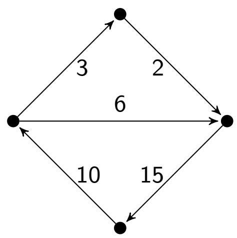

Figure 1: An oriented $\mathbb{N}$-labelled graph (G, l)

We assign to each of its three circuits the clockwise travelling direction. We obtain a circuit matrix of G and a labelled circuit matrix of (G, l):

$$M_G = \begin{bmatrix} 1 & 1 & -1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 \\ 1 & 1 & 0 & 1 & 1 \end{bmatrix} \qquad M_{(G,l)} = \begin{bmatrix} 3 & 2 & -6 & 0 & 0 \\ 0 & 0 & 6 & 15 & 10 \\ 3 & 2 & 0 & 15 & 10 \end{bmatrix}$$

Choose the spanning tree with edges labelled by $3, 6$ and 10. The fundamental circuit matrix of G and lfc-matrix of (G, l) are obtained from M_G and $M_{(G,l)}$ by removing the third row:

$$N_G = \begin{bmatrix} 1 & 1 & -1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 \end{bmatrix} \qquad N_{(G,l)} = \begin{bmatrix} 3 & 2 & -6 & 0 & 0 \\ 0 & 0 & 6 & 15 & 10 \end{bmatrix}$$

Let M be an integer-valued matrix with a rows and b columns. There exist matrices $A \in \mathrm{GL}(a, \mathbb{Z})$ and $B \in \mathrm{GL}(b, \mathbb{Z})$ such that

$$AMB = \begin{bmatrix} d_1 & 0 & 0 & & \dots & & 0 \\ 0 & d_2 & 0 & & \dots & & 0 \\ 0 & 0 & \ddots & & & & 0 \\ \vdots & & & d_k & & & \vdots \\ & & & & 0 & & \\ & & & & & \ddots & \\ 0 & & & \dots & & & 0 \end{bmatrix}$$

where the diagonal entries satisfy $d_i | d_{i+1}$ for $i = 1, \dots, k-1$. This is the so-called *Smith normal form* of M and it is unique up to multiplication of the diagonal entries by units of $\mathbb{Z}$. For $1 \leq i \leq k$, the integer d_i is the quotient D_i/D_{i-1}, where D_i equals the greatest common divisor of all minors of order i of M.

Going back to the matrices $M_{(G,l)}$ and its submatrix $N_{(G,l)}$, it follows from theorem 11.3 that their Smith normal forms both have rank equal to the nullity r of the graph G. Besides, as any row of $M_{(G,l)}$ is a $\mathbb{Z}$-linear combination of rows of $N_{(G,l)}$, we see that the numbers D_i defined above are the same for the two matrices. It follows that $M_{(G,l)}$ and $N_{(G,l)}$ have the same non-zero numbers d_i appearing on the diagonal. Moreover, the numbers $d_1, \ldots, d_r$ are defined up to multiplication by -1, hence do not depend on the choices of orientation of edges or circuits, but only on the $\mathbb{N}$-labelled graph (G, l).

11.3 Cartier labellings and blow-up graphs

Let (G, l) be an $\mathbb{N}$-labelled graph. Let $\mathbb{Z}^V$ be the free abelian group generated by the set of vertices V. Any element φ of $\mathbb{Z}^V$ can be interpreted as a vertex labelling $\varphi\colon V \to \mathbb{Z}$ of the graph G.

Definition 11.6. An element $\varphi \in \mathbb{Z}^V$ is a *Cartier vertex labelling* if for every edge $e \in E$ with endpoints $v, w \in V$, $l(e)$ divides $\varphi(v) - \varphi(w)$.

We denote by $\mathcal{C} \subset \mathbb{Z}^V$ the subgroup of Cartier vertex labellings.

Definition 11.7. We call *multidegree operator* the group homomorphism $\delta\colon \mathcal{C} \to \mathbb{Z}^V$ which sends $\varphi \in \mathcal{C}$ to

$$v \mapsto \sum_{\substack{\text{edges } e \\ \text{incident to } v}} \frac{\varphi(w) - \varphi(v)}{l(e)}$$

where w denotes the other endpoint of e (which is v itself if e is a loop).

Lemma 11.8. *The kernel of δ consists of the constant vertex labellings, hence there is an exact sequence*

$$0 \to \mathbb{Z} \to \mathcal{C} \xrightarrow{\delta} \mathbb{Z}^V.$$

Proof. Any constant vertex labelling is in the kernel of δ. Conversely, let $\varphi \in \ker \delta$ and let $v \in V$ be a vertex where φ attains its maximum. Then for all the vertices w adjacent to v one has $\varphi(w) = \varphi(v)$. Since the graph is finite and connected, one can repeat the argument and find that φ is a constant labelling. $\square$

Remark 11.9. When the edge-labelling $l\colon E \to \mathbb{Z}_{\geq 1}$ is constant with value 1, the multidegree operator δ coincides with the Laplacian operator of the graph G.

Definition 11.10. Given an $\mathbb{N}$-labelled graph $(G, l) = (V, E, l)$ we define the *total blow-up graph* $(\widetilde{G}, \widetilde{l}) = (\widetilde{V}, \widetilde{E}, \widetilde{l})$ to be the $\mathbb{N}$-labelled graph constructed as follows starting from (G, l): every edge $e \in E$ is replaced by a path consisting of $l(e)$ edges, and $\widetilde{l}\colon \widetilde{E} \to \mathbb{Z}$ is set to be the constant labelling with value 1.

Example 11.11. Figure 2 shows an $\mathbb{N}$-labelled graph (a) and its total blow-up graph (b).

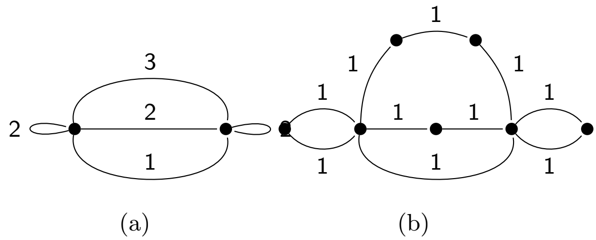

Figure 2: An $\mathbb{N}$-labelled graph G (a) and its total blow-up graph $\widetilde{G}$ (b).

We call *old vertices* the vertices in the image of the inclusion map $V \hookrightarrow \widetilde{V}$. We call *new vertices* the remaining vertices.

Notice that every new vertex is incident to exactly two edges, and belongs to a unique path (corresponding to some edge $e \in E$) connecting two old vertices of $\widetilde{V}$. Just as before we consider the group of Cartier vertex labellings $\widetilde{\mathcal{C}}$ of $(\widetilde{G}, \widetilde{l})$, and the multidegree operator $\widetilde{\delta}\colon \widetilde{\mathcal{C}} \to \mathbb{Z}^{\widetilde{V}}$.

We obtain a morphism of exact sequences

$$\begin{array}{ccccccc} 0 & \longrightarrow & \mathbb{Z} & \longrightarrow & \mathcal{C} & \xrightarrow{\delta} & \mathbb{Z}^V \\ & & \downarrow{\scriptstyle \mathrm{id}} & & \downarrow{\scriptstyle \iota} & & \downarrow{\scriptstyle \epsilon} \\ 0 & \longrightarrow & \mathbb{Z} & \longrightarrow & \widetilde{\mathcal{C}} & \xrightarrow{\widetilde{\delta}} & \mathbb{Z}^{\widetilde{V}} \end{array} \tag{30}$$

The map $\epsilon\colon \mathbb{Z}^V \to \mathbb{Z}^{\widetilde{V}}$ is given by extending vertex-labellings by zero on the set of new vertices. The map $\iota\colon \mathcal{C} \to \widetilde{\mathcal{C}}$ sends a Cartier vertex labelling φ on G to the Cartier vertex labelling $\iota(\varphi)$ on $\widetilde{G}$ whose value at old vertices in inherited by φ, and extended by linear interpolation to the new vertices.

More precisely: if e is an edge of G with endpoints v, w which is replaced in $\widetilde{G}$ by a path consisting of vertices $v = v_0, v_1, \ldots, v_{l(e)} = w$, we set for each $k = 0, \ldots, l(e)$

$$\iota(\varphi)(v_k) = \frac{(l(e) - k)\varphi(v) + k\varphi(w)}{l(e)}.$$

The Cartier condition on φ implies that this labelling takes integer values.

Let $H = \operatorname{coker} \delta$, $\widetilde{H} = \operatorname{coker} \widetilde{\delta}$. The commutative diagram above yields a group homomorphism $\overline{\epsilon}\colon H \to \widetilde{H}$.

Lemma 11.12. *The group homomorphism $\overline{\epsilon}\colon H \to \widetilde{H}$ is injective.*

Proof. Let $\alpha \in \mathbb{Z}^V$ be a vertex labelling and let $\epsilon(\alpha) \in \mathbb{Z}^{\widetilde{V}}$ be its extension by zero. Assume that there exists a Cartier vertex labelling $\widetilde{\varphi} \in \widetilde{\mathcal{C}}$ such that $\epsilon(\alpha) = \widetilde{\delta}(\widetilde{\varphi})$. Then $\widetilde{\delta}(\widetilde{\varphi})$ takes value zero on all new vertices of $\widetilde{G}$. Hence, if v is a new vertex of $\widetilde{G}$ adjacent to two verteces v' and v'', we have $\widetilde{\varphi}(v') - \widetilde{\varphi}(v) = \widetilde{\varphi}(v) - \widetilde{\varphi}(v'')$. We immediately see that $\widetilde{\varphi}$ is an interpolation of a Cartier vertex labelling $\varphi \in \mathcal{C}$, i.e. $\widetilde{\varphi}$ is in the image of ι. Since $\epsilon\colon \mathbb{Z}^V \to \mathbb{Z}^{\widetilde{V}}$ is injective, $\alpha = \delta(\varphi)$. □

Our aim now is to give necessary and sufficient conditions on the $\mathbb{N}$-labelled graph (G, l) for the map $\overline{\epsilon}\colon H \to \widetilde{H}$ to be surjective (hence an isomorphism).

11.4 Circuit-coprime graphs

Definition 11.13. Let $(G, l) = (V, E, l)$ be an $\mathbb{N}$-labelled graph. We say that (G, l) is *circuit-coprime* if for every circuit $C \subset G$, $\gcd\{l(e) | e \text{ is an edge of } C\} = 1$.

Example 11.14. In fig. 3 the $\mathbb{N}$-labelled graph (a) is circuit-coprime, whereas the $\mathbb{N}$ -labelled graph (b) is not, as it contains a loop labelled by 3 in addition to a circuit labelled by 6, 10 and 10.

Lemma 11.15. *Let $(G, l) = (V, E, l)$ be an $\mathbb{N}$-labelled graph. Denote by r its nullity. The Smith normal form of the matrix $M_{(G,l)}$ has diagonal entries $d_1 = d_2 = \ldots = d_r = 1$ if and only if (G, l) is circuit-coprime.*

Proof. Assume first that (G, l) is not circuit-coprime. Let C be a circuit whose labels have greatest common divisor $D \neq 1$. Pick an edge e of C. The subgraph $C \setminus e$ is a tree; let T be a spanning tree of G containing it. Then e is a link for T, and C is its associated fundamental circuit. The lfc-matrix $N_{(G,l)}$ has a row corresponding to the circuit C, hence all entries of this row are divisible

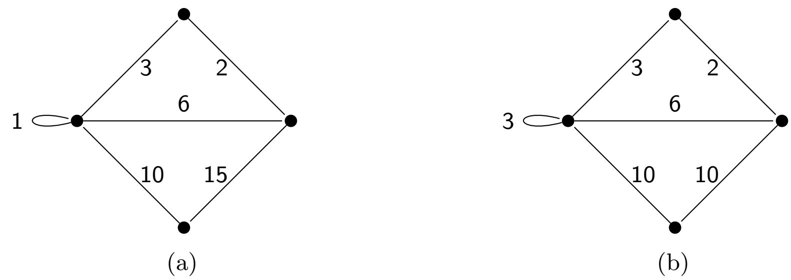

Figure 3: A circuit-coprime $\mathbb{N}$-labelled graph (a) and an $\mathbb{N}$ -labelled graph that is not circuit-coprime (b).

by D. Then the linear map $f\colon \mathbb{Z}^n \to \mathbb{Z}^r$ defined by $N_{(G,l)}$ is not surjective; hence the linear map associated to the Smith normal form of $N_{(G,l)}$ is not surjective either. Therefore, some (necessarily non-zero) diagonal entry of the Smith normal form of $N_{(G,l)}$ is different from ± 1. As previously remarked, the Smith normal forms of $M_{(G,l)}$ and $N_{(G,l)}$ have the same non-zero diagonal entries, hence $d_r \neq \pm 1$.

Conversely, assume that G is circuit-coprime. After fixing some spanning tree T, consider the lfc-matrix $N_{(G,l)}$. We only need to prove that the diagonal entries of the Smith normal form of $N_{(G,l)}$ are all 1, which amounts to proving that the greatest common divisor d of the minors of order r of the lfc-matrix $N_{(G,l)}$ is 1.

As we have seen in section 11.2, we have the relation

$$N_{(G,l)} = N_G \cdot L.$$

Let N' be a maximal square submatrix of $N_{(G,l)}$. Then N' corresponds to r edges of G, which we denote $e_{i_1}, e_{i_2}, \ldots, e_{i_r}$. Let N'' be the corresponding square submatrix of N_G. We have the relation

$$\det N' = \prod_{j=1}^{r} l(e_{i_j}) \det N''$$

By [TS92], Theorem 6.15, all minors of N_G are either $1, 0$ or -1, hence $\det N''$ is either $1, 0$ or -1. Moreover, by [TS92], Theorem 6.10, a square submatrix of order r of N_G has determinant ± 1 if and only if the corresponding r edges are the complement of a spanning tree. Hence $\det N' = \pm \prod_{i=1}^{r} l(e_{i_j})$ if the edges $e_{i_1}, e_{i_2}, \ldots, e_{i_r}$ form the complement of a spanning tree of G, otherwise $\det N' = 0$. We claim that

$$d := \gcd\{\det N' | N' \text{ is an } r \times r \text{ square submatrix of } N_{(G,l)}\} = 1.$$

Let p be a prime number and denote by E_p the set of edges e of G whose label $l(e)$ is divisible by p. Because (G, l) is circuit-coprime, E_p contains no circuit; hence E_p is contained in some spanning tree T of G. There are exactly r edges, $e_1, e_2, \ldots, e_r$, that do not belong to T. These give a square $r \times r$ submatrix of $N_{(G,l)}$ whose determinant is $\prod_{i=1}^{r} l(e_i) \not\equiv 0 \pmod p$, since $e_1, \ldots, e_r \notin E_p$. Hence $p \nmid d$. It follows that $d = 1$; since $d_i | d_{i+1}$ for all $i = 1, \ldots, r-1$ and $d_r | d$, we obtain the result. □

Proposition 11.16. *Let $(G, l) = (V, E, l)$ be an $\mathbb{N}$-labelled graph. The group homomorphism $\overline{\epsilon}\colon H \to \widetilde{H}$ is an isomorphism if and only if (G, l) is circuit-coprime.*

Proof. We already know that $\overline{\epsilon}\colon H \to \widetilde{H}$ is injective by lemma 11.12. It is surjective if and only if for every vertex-labelling $\alpha \in \mathbb{Z}^{\widetilde{V}}$, there exists $\widetilde{\varphi} \in \widetilde{\mathcal{C}}$ such that $\widetilde{\delta}(\widetilde{\varphi}) + \alpha$ is in the image of the extension-by-zero map $\epsilon\colon \mathbb{Z}^V \to \mathbb{Z}^{\widetilde{V}}$, i.e. $\widetilde{\delta}(\widetilde{\varphi}) + \alpha$ is supported on the set of old vertices. We may of course assume that α belongs to the canonical basis of $\mathbb{Z}^{\widetilde{V}}$. That is, $\alpha = \chi_v$ for some vertex v of $\widetilde{G}$, where

$$\chi_v(w) = \begin{cases} 1 & \text{if } w = v \\ 0 & \text{if } w \neq v. \end{cases}$$

If v is an old vertex of $\widetilde{G}$, χ_v is an extension by zero of a vertex-labelling on G, so we may assume that v is a new vertex. Then v belongs to some path $P \subset \widetilde{G}$ associated to some edge $\overline{e} \in E$. Denote by $w_0, w_1, \ldots, w_{l(\overline{e})}$ the vertices of the path P, so that w_0 and $w_{l(\overline{e})}$ are old vertices, and the numbering of the indices follows the order of the vertices on the path. For every $i = 1, \ldots, l(\overline{e})$, let $\alpha_i = \chi_{w_i} - \chi_{w_0} \in \mathbb{Z}^{\widetilde{V}}$ be the vertex-labelling that has value 1 at w_i, value -1 at w_0, and value 0 everywhere else. Then it is easy to check that the images $\overline{\alpha}_i$ of the α_i in $\widetilde{H}$ satisfy $k\overline{\alpha}_1 = \overline{\alpha}_k$ for all $k = 1, \ldots, l(\overline{e})$. Hence, if $\overline{\alpha_1}$ is in the image of $\overline{\epsilon}\colon H \to \widetilde{H}$, so are all the $\overline{\alpha_i}$ for $1 \leq i \leq l(e)$. This shows that we can take v to be equal to w_1; hence $\chi_v = \chi_{w_1}$ takes value 1 on a new vertex v adjacent to an old vertex, and value zero at all other vertices.

We ask whether an element $\widetilde{\varphi} \in \widetilde{\mathcal{C}}$ exists such that $\widetilde{\delta}(\widetilde{\varphi}) + \chi_{w_1}$ is supported only on the old vertices. In other words, $\widetilde{\delta}(\widetilde{\varphi})$ must be zero on all new vertices except for the vertex w_1, where it has to take the value -1. This is equivalent to asking that, for every new vertex z, adjacent to vertices z_1, z_2,

$$\begin{cases} \widetilde{\varphi}(z) - \widetilde{\varphi}(z_1) = \widetilde{\varphi}(z_2) - \widetilde{\varphi}(z) & \text{if } z \neq w_1 \\ (\widetilde{\varphi}(z_1) - \widetilde{\varphi}(z)) + (\widetilde{\varphi}(z_2) - \widetilde{\varphi}(z)) = -1 & \text{if } z = w_1 \end{cases} \tag{31}$$

holds.

We claim that such a $\widetilde{\varphi}$ exists if and only if there exists a vertex-labelling φ of the graph G, such that, for every edge $e \in E$ with endpoints v_0, v_1,

$$\begin{cases} \varphi(v_1) - \varphi(v_0) \equiv 0 \bmod l(e) & \text{if } e \neq \overline{e} \\ \varphi(v_1) - \varphi(v_0) \equiv 1 \bmod l(e) & \text{if } e = \overline{e}, v_0 = w_0, v_1 = w_{l(e)} \end{cases} \tag{32}$$

where we have identified the old vertices $w_0, w_{l(e)}$ with the corresponding vertices in G. Indeed, given $\widetilde{\varphi}$ one obtains φ simply by restriction to old vertices. Conversely, given a φ as in (37), $\widetilde{\varphi}$ is obtained as follows: for an edge $e \neq \overline{e}$, we define $\widetilde{\varphi}$ on the corresponding path $\{z_0 = v_0, z_1, z_2, \ldots, z_{l(e)} = v_1\}$ by:

$$\forall k = 0, 1, \ldots, l(e), \quad \widetilde{\varphi}(z_k) = \frac{k\varphi(v_1) + (l(e) - k)\varphi(v_0)}{l(e)}.$$

On the path $\{w_0 = v_0, w_1, \ldots, w_{l(\overline{e})} = v_1\}$ corresponding to the edge $\overline{e}$, we set instead

$$\widetilde{\varphi}(w_k) = \begin{cases} \frac{k\varphi(v_1) + (l(e) - k)(\varphi(v_0) + 1)}{l(e)} & \text{if } k \in \{1, 2, \ldots, l(\overline{e})\} \\ \widetilde{\varphi}(v_0) & \text{if } k = 0; \end{cases}$$

which establishes the claim.

If the graph G is a tree it is clear that such a φ can be found. If there are circuits in G, the existence of a solution φ depends of course on the labels of the circuits. Fix an orientation on G, so that we have source and target functions $s, t \colon E \to V$, and so that $s(\overline{e}) = w_0$, $t(\overline{e}) = w_{l(\overline{e})}$. Assume that a vertex-labelling φ of G satisfying the conditions (37) exists. In particular we have that $\varphi(t(\overline{e})) - \varphi(s(\overline{e})) \equiv 1 \bmod l(\overline{e})$. For every edge $e \in E$ let

$$x(e) := \begin{cases} \frac{\varphi(t(e)) - \varphi(s(e))}{l(e)} & \text{if } e \neq \overline{e} \\ \frac{\varphi(t(e)) - \varphi(s(e)) - 1}{l(e)} & \text{if } e = \overline{e} \end{cases}$$

Let $C \subset G$ be a circuit consisting of vertices $v_0, v_1, \ldots, v_s = v_0$ connected by edges $e_0, e_1, e_2, \ldots, e_s = e_0$, so that e_i connects v_i and v_{i+1} for every $i \in \mathbb{Z}/s\mathbb{Z}$. Notice that the increasing numbering gives an orientation to C. We have

$$(\varphi(v_s) - \varphi(v_{s-1})) + (\varphi(v_{s-1}) - \varphi(v_{s-2})) + \ldots + (\varphi(v_1) - \varphi(v_s)) = 0.$$

Setting

$$a_i = \begin{cases} 1 & \text{if } t(e_i) = v_{i+1}, s(e_i) = v_i \\ -1 & \text{if } t(e_i) = v_i, s(e_i) = v_{i+1} \end{cases} \tag{33}$$

for every $i \in \mathbb{Z}/s\mathbb{Z}$, we obtain

$$\sum a_i x_{e_i} l(e_i) = 0$$

if the edge $\overline{e}$ does not belong to the circuit C, whereas if $\overline{e} \in C$ we have

$$\sum a_i x_{e_i} l(e_i) = \begin{cases} -1 & \text{if the orientations of } C \text{ and } \overline{e} \text{ agree;} \\ 1 & \text{if the orientations of } C \text{ and } \overline{e} \text{ do not agree;} \end{cases}$$

Let $C_1, \dots, C_m$ be the circuits of G. Choose an orientation for each circuit, so that we can form the labelled circuit matrix $M_{(G,l)}$ associated to G. We see that the vector $\underline{x} = (x_1, \dots, x_n)$ is a solution of

$$M_{(G,l)} x = b(\overline{e})$$

where $b(\overline{e}) = (b_1, \dots, b_m)$ with

$$b_i = \begin{cases} 0 & \text{if } \overline{e} \notin C_i; \\ -1 & \text{if } \overline{e} \in C_i \text{ and the orientation of } \overline{e} \text{ agrees with the} \\ & \text{orientation of } C_i; \\ 1 & \text{if } \overline{e} \in C_i \text{ and the orientation of } \overline{e} \text{ does not agree with the} \\ & \text{orientation of } C_i. \end{cases}$$

Conversely, a solution $x \in \mathbb{Z}^n$ to the system $M_{(G,l)} x = b(\overline{e})$ yields a vertex labelling φ as in (37). We conclude that the map $\overline{\epsilon}\colon H \to \widetilde{H}$ is surjective if and only if for every edge $e \in E$, there is a solution $x \in \mathbb{Z}^n$ to

$$M_{(G,l)} x = b(e).$$

After having chosen a spanning tree T and formed the lfc-matrix $N_{(G,l)}$, this is in turn equivalent to the map $\mathbb{Z}^n \to \mathbb{Z}^r$ defined by $N_{(G,l)}$ being surjective. Indeed, the set $\{b(e) | e \text{ is a link of } T\}$ is a basis for $\mathbb{Z}^r$. Now, $N_{(G,l)}$ is surjective if and only if its Smith normal form (or equivalently the one of $M_{(G,l)}$) has only 1's on the diagonal. By lemma 11.15, we conclude.

□

11.5 $\mathbb{N}_\infty$-labelled graphs

We want to generalize the results of the previous subsection to labelled graphs whose labels can attain the value ∞. Denote by $\mathbb{N}_\infty$ the set $\mathbb{Z}_{\geq 1} \cup \{\infty\}$. Let $(G, l) = (V, E, l)$ be the datum of a graph, with set of vertices V and set of edges E, and of a function $l\colon E \to \mathbb{N}_\infty$. We say that (G, l) is an $\mathbb{N}_\infty$-*labelled graph.*

The notions of Cartier vertex labelling 11.6 and multidegree operator 11.7 carry over to this setting without change, imposing that the only integer divisible by ∞ is 0, and setting $\frac{0}{\infty} = 0$ in the definition of multidegree operator.

In particular, if a vertex-labelling on (G, l) is Cartier, it attains the same value at the two extremal vertices of an edge with label ∞.

Definition 11.17. Given an $\mathbb{N}_\infty$-labelled graph $(G, l) = (V, E, l)$ we define the *first-blow-up graph* $G_1 = (V_1, E_1, l_1)$ to be the $\mathbb{N}_\infty$-labelled graph constructed as follows starting from (G, l): every edge $e \in E$ with $l(e) = 1$ is preserved unaltered; every edge $e \in E$ with $l(e) \geq 2$ is replaced by a path consisting of an edge labelled by 1, followed by an edge labelled by $l(e) - 2$ (which could equal 0 or ∞), followed by an edge labelled by 1.

We define inductively for every integer $n \geq 1$ the *n-th blow-up graph* $G_n = (V_n, E_n, l_n)$ as the first-blow-up graph of G_{n-1}.

Example 11.18. Figure 4 shows an $\mathbb{N}_\infty$-labelled graph (a) with its first (b) and second (c) blow-up graphs.

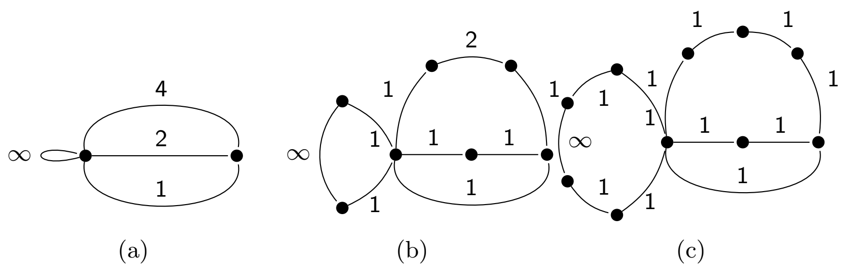

Figure 4: An $\mathbb{N}_\infty$-labelled graph (a) with its first (b) and second (c) blow-up graphs

Denote by $\mathcal{C}_n$ the group of Cartier vertex-labellings on (G_n, l_n). Just as in

(30), we obtain a commutative diagram

$$\begin{array}{ccc}
\mathcal{C} & \xrightarrow{\delta} & \mathbb{Z}^{V} \\
\downarrow{\scriptstyle \iota_1} & & \downarrow{\scriptstyle \epsilon_1} \\
\mathcal{C}_1 & \xrightarrow{\delta_1} & \mathbb{Z}^{V_1} \\
\downarrow{\scriptstyle \iota_2} & & \downarrow{\scriptstyle \epsilon_2} \\
\vdots & & \vdots \\
\downarrow & & \downarrow \\
\mathcal{C}_n & \xrightarrow{\delta_n} & \mathbb{Z}^{V_n} \\
\downarrow{\scriptstyle \iota_n} & & \downarrow{\scriptstyle \epsilon_n} \\
\vdots & & \vdots
\end{array}$$

The vertical maps ϵ_j are once again extension by zero; the maps ι_j are defined as follows: if e is an edge of G_{j-1} which is replaced in G_j by a path consisting of vertices $v_0 = v, v_1, v_2, v_3 = w$ (with possibly $v_1 = v_2$, if $l_{j-1}(e) = 2$), and φ is Cartier vertex labelling on G_{j-1}, we set $\iota_j(\varphi)$ to take the value $\varphi(v)$ at v_0, $\frac{(l(e)-1)\varphi(v)+\varphi(w)}{l(e)}$ at v_1, $\frac{\varphi(v)+(l(e)-1)\varphi(w)}{l(e)}$ at v_2, $\varphi(w)$ at v_3. The diagram above gives rise to a chain of group homomorphisms

$$H \to H_1 \to H_2 \to \ldots \to H_n \to \ldots \tag{34}$$

between the cokernels of the rows. Each map of the chain (34) is injective; we ask whether they are all isomorphisms, i.e. under which conditions

$$H \to \operatorname{colim} H_i \tag{35}$$

is an isomorphism.

Definition 11.19. Let $(G, l) = (V, E, l)$ be an $\mathbb{N}_\infty$-labelled graph. We let $(G, l^\circ) = (V, E, l^\circ)$ be the $\mathbb{N}_0 := \mathbb{Z}_{\geq 0}$-labelled graph obtained from (G, l) by setting $l^\circ(e) = 0$ for all edges e with label $l(e) = \infty$.

We say that (G, l) is *circuit-coprime* if for every circuit $C \subset G$,

$$\gcd(l^\circ(e) | e \text{ is an edge of } C) = 1.$$

Here we define the gcd of a subset $S \subset \mathbb{Z}$ to be the non-negative generator of the ideal $\langle S \rangle \subset \mathbb{Z}$.

Remark 11.20. An $\mathbb{N}_\infty$-labelled graph containing a circuit whose labels are all ∞ is not circuit-coprime. Indeed, $\gcd(0) = 0$.

Proposition 11.21. *Let (G, l) be an $\mathbb{N}_\infty$-labelled graph. The map* (35) *is an isomorphism if and only if (G, l) is circuit-coprime.*

Proof. Instead of (G, l) and its blow-up graphs $(G_1, l_1), (G_2, l_2), \dots$ we consider $(G, l^\circ), (G_1, l_1^\circ), (G_2, l_2^\circ), \dots$. We keep the same notion of Cartier vertex labelling and multidegree operator, by imposing that the only integer divisible by 0 is 0, and that $0/0 = 0$. The chain of homomorphisms 34 is also preserved. To keep the notation light, we drop the $^\circ$'s. From now on, the proof is a readaptation of the content of section 11.4. First, for labelled graphs whose labels attain the value 0, we define the labelled circuit matrix $M_{(G,l)}$ and labelled fundamental circuit matrix $N_{(G,l)}$ in the same way as in section 11.2. Lemma 11.15 stays true in this setting, so we find that (G, l) is circuit-coprime if and only if $N_{(G,l)}$ is surjective.

To finish the proof we only need to readapt proposition 11.16 to our new setting. So, we want to show that $N_{(G,l)}$ is surjective if and only if $\epsilon_n : H \to H_n$ is surjective for all $n \geq 1$. We fix an integer n big enough, so that all labels of G_n are 1's or 0's. As in proposition 11.16, we let $\alpha \in \mathbb{Z}^{V_n}$; we may pick $\alpha = \chi_v$ for some vertex v belonging to some path $P \subset G_n$ associated to some edge $\overline{e} \in E$. Denote by $w_0, w_1, \dots, w_r$ the vertices of the path P. We may still assume that $v = w_1$. Indeed, if there is no edge in P labelled by zero, one reasons as in proposition 11.16; otherwise, if there is an edge in P labelled by 0, then it has to be the edge connecting w_s and w_{s+1}, with $s = \frac{r-1}{2}$. We may assume without loss of generality that $v = w_k$ for $k \leq s$. We get that $\overline{\chi_{w_k}} = k\overline{\chi_{w_1}}$ in H_n (as always, compare with proposition 11.16).

An element $\tilde{\varphi}$ in $\mathcal{C}_n$ is such that $\delta_n(\varphi_n) + \chi_{w_1}$ is supported on the old vertices is a vertex-labelling $\tilde{\varphi} \in Z^{V_n}$ satisfying the following: for every new vertex z, adjacent to vertices z_1 and z_2,

$$\begin{cases} \widetilde{\varphi}(z) - \widetilde{\varphi}(z_1) = \widetilde{\varphi}(z_2) - \widetilde{\varphi}(z) & \text{if } z \neq w_1 \\ (\widetilde{\varphi}(z_1) - \widetilde{\varphi}(z)) + (\widetilde{\varphi}(z_2) - \widetilde{\varphi}(z)) = -1 & \text{if } z = w_1 \end{cases} \tag{36}$$

That such a Cartier vertex-labelling exists means that there is a vertex-labelling $\tilde{\varphi}$ satisfying the condition 36 above, plus the extra condition that $\tilde{\varphi}(z_1) = \tilde{\varphi}(z_2)$ for any two adjacent vertices z_1, z_2 connected by an edge labelled by zero.

In turn, such a $\widetilde{\varphi}$ exists if and only if there exists a vertex-labelling φ of G such that, for every edge $e \in E$ with endpoints v_0, v_1,

$$\begin{cases} \varphi(v_1) - \varphi(v_0) \equiv 0 \bmod l(e) & \text{if } e \neq \overline{e} \\ \varphi(v_1) - \varphi(v_0) \equiv 1 \bmod l(e) & \text{if } e = \overline{e}, v_0 = w_0, v_1 = w_r \end{cases} \tag{37}$$

where we have identified the old vertices w_0, w_r with the corresponding vertices in G. This is the same condition as condition 37 in proposition 11.16. From this point on, the rest of the proof coincides with the proof of proposition 11.16; we only mention that, at the point when $x(e)$ is defined, one can assign to it any value if $l(e) = 0$.

□

12 Semi-factoriality of nodal curves

Let S be the spectrum of a discrete valuation ring R having perfect fraction field K, residue field k and uniformizer t. Let $f\colon \mathcal{X} \to S$ be a nodal curve whose special fibre has split singularities, and $\Gamma = (V, E)$ be the dual graph of the special fibre $\mathcal{X}_k$. For any $v \in V$, we denote by X_v the corresponding irreducible component of the special fibre $\mathcal{X}_k$.

Definition 12.1. The *labelled graph* of $\mathcal{X} \to S$ is the $\mathbb{N}_\infty$-labelled graph (Γ, l) whose labelling l assigns to each edge of Γ the thickness (see section 7.1) of the corresponding singular point of $\mathcal{X}_k$.

Our aim is to relate the property of being circuit-coprime for the graph (Γ, l) to the semi-factoriality of $f\colon \mathcal{X} \to S$. To this end, we are going to provide a dictionary between the geometry of $\mathcal{X}/S$ and the combinatorial objects introduced in section 11.

Denote by $\mathrm{Div}_k(\mathcal{X})$ the group of Weil divisors on $\mathcal{X}$ supported on the special fibre $\mathcal{X}_k$. It is the free abelian group generated by the irreducible components of $\mathcal{X}_k$. Hence we obtain a natural isomorphism $\mathrm{Div}_k(\mathcal{X}) \to \mathbb{Z}^V$.

Let $\mathcal{C}(\mathcal{X})$ be the group of Cartier divisors on $\mathcal{X}$ whose restriction to the generic fibre $\mathcal{X}_K$ is trivial. We claim that the natural map $\mathcal{C}(\mathcal{X}) \to \mathrm{Div}_k(\mathcal{X})$ is injective. This follows from ([GD67], 21.6.9 (i)) under the assumption that $\mathcal{X}$ is normal, which is not satisfied if $\mathcal{X}/S$ has singular generic fibre. However, the proof only requires that for all $x \in \mathcal{X}_k$, $\mathrm{depth}(\mathcal{O}_{\mathcal{X},x}) = 1$ implies $\dim \mathcal{O}_{\mathcal{X},x} = 1$. This is immediately checked: let $x \in \mathcal{X}_k$ with $\dim \mathcal{O}_{\mathcal{X},x} \neq 1$; then x is a closed point of $\mathcal{X}_k$. By S-flatness of $\mathcal{X}$, the uniformizer t is not a zero divisor in $\mathcal{O}_{\mathcal{X},x}$; as $\mathcal{X}_k$ is reduced, $\mathcal{O}_{\mathcal{X},x}/t\mathcal{O}_{\mathcal{X},x}$ is reduced. Every reduced noetherian ring of dimension 1 is Cohen-Macaulay, hence $\mathrm{depth}(\mathcal{O}_{\mathcal{X},x}/t\mathcal{O}_{\mathcal{X},x}) = 1$, and we deduce by [Sta16]TAG 0AUI that $\mathrm{depth}(\mathcal{O}_{\mathcal{X},x}) = 2$, establishing the claim. Hence $\mathcal{C}(\mathcal{X})$ is in a natural way a subgroup of $\mathrm{Div}_k(\mathcal{X})$.

Finally, denote by $E(\mathcal{X})$ the kernel of the restriction map $\mathrm{Pic}(\mathcal{X}) \to \mathrm{Pic}(\mathcal{X}_K)$, so that $E(\mathcal{X})$ is the group of isomorphism classes of line bundles on $\mathcal{X}$ that are generically trivial. We have an exact sequence of groups

$$0 \to \mathbb{Z} \to \mathcal{C}(\mathcal{X}) \to E(\mathcal{X}) \to 0$$