Around the numeric-symbolic computation of differential Galois groups

Joris van der Hoeven

Dépt. de Mathématiques (bât. 425)
Université Paris-Sud
91405 Orsay Cedex
France

Email: `vdhoeven@texmacs.org`

July 15, 2006

Let $L \in \mathbb{K}(z)[\partial]$ be a linear differential operator, where $\mathbb{K}$ is an effective algebraically closed subfield of $\mathbb{C}$. It can be shown that the differential Galois group of L is generated (as a closed algebraic group) by a finite number of monodromy matrices, Stokes matrices and matrices in local exponential groups. Moreover, there exist fast algorithms for the approximation of the entries of these matrices.

In this paper, we present a numeric-symbolic algorithm for the computation of the closed algebraic subgroup generated by a finite number of invertible matrices. Using the above results, this yields an algorithm for the computation of differential Galois groups, when computing with a sufficient precision.

Even though there is no straightforward way to find a "sufficient precision" for guaranteeing the correctness of the end-result, it is often possible to check *a posteriori* whether the end-result is correct. In particular, we present a non-heuristic algorithm for the factorization of linear differential operators.

1. Introduction

Let $L \in \mathbb{K}(z)[\partial]$ be a monic linear differential operator of order n, where $\mathbb{K}$ is an effective algebraically closed subfield of $\mathbb{C}$. A holonomic function is a solution to the equation $L f = 0$. The differential Galois group $\mathcal{G}$ of L is a linear algebraic group which acts on the space H of solutions (see section 2.2 and [Kap57, vdPS03, Kol73]). It carries a lot of information about the solutions in H and on the relations between different solutions. For instance, the existence of non-trivial factorizations of L and the existence of Liouvillian solutions can be read off from the Galois group. This makes it an interesting problem to explicitly compute the Galois group of L.

A classical approach in this area is to let $\mathcal{G}$ act on other vector spaces obtained from H by the constructions from linear algebra, such as symmetric powers $\otimes^k H$ and exterior powers $\wedge^k H$ [Bek94, SU93]. For a suitable such space S, the Galois group $\mathcal{G}$ consists precisely of those invertible $n \times n$ matrices which leave a certain one-dimensional subspace of S invariant [Hum81, chapter 11]. Invariants in $\otimes^k H$ or $\wedge^k H$ under $\mathcal{G}$ may be computed more efficiently by considering the local solutions of $L f = 0$ at singularities [vHW97, vH97, vH96]. More recently, and assuming (for instance) that the coefficients of L are actually in $\mathbb{Q}(z)$, alternative algorithms appeared which are based on the reduction of the equation $L f = 0$ modulo a prime number p [Clu04, vdP95, vdPS03].

In this paper, we will study another type of "analytic modular" algorithms, by studying the operator L in greater detail near its singularities using the theory of accelero-summation [É85, É87, É92, É93, Bra91, Bra92]. More precisely, we will use the following facts:

- The differential Galois group of L is generated (as a closed algebraic group) by a finite number of monodromy matrices, Stokes matrices and matrices in so called local exponential groups.

- There exists an algorithm [vdH99, vdH01b, vdH05b] for the approximation of the entries of the above matrices. If $\mathbb{K}=\mathbb{Q}^{\mathrm{alg}}$ is the algebraic closure of $\mathbb{Q}$, then d-digit approximations can be computed in time $O(d\log^4 d\log\log d)$.

When using these facts for the computation of differential Galois groups, the bulk of the computations is reduced to linear algebra in dimension n with multiple precision coefficients.

In comparison with previous methods, this approach is expected to be much faster than algorithms which rely on the use of exterior powers. A detailed comparison with arithmetic modular methods would be interesting. One advantage of arithmetic methods is that they are easier to implement in existing systems. On the other hand, our analytic approach relies on linear algebra in dimension n (with floating coefficients), whereas modulo p methods rely on linear algebra in dimension $n\,p$ (with coefficients modulo p), so the first approach might be a bit faster. Another advantage of the analytic approach is that it is more easily adapted to coefficients fields $\mathbb{K}$ with transcendental constants.

Let us outline the structure of this paper. In section 2, we start by recalling some standard terminology and we shortly review the theorems on which our algorithms rely. We start with a survey of differential Galois theory, monodromy and local exponential groups. We next recall some basic definitions and theorems from the theory of accelero-summation and the link with Stokes matrices and differential Galois groups. We finally recall some theorems about the effective approximation of the transcendental numbers involved in the whole process.

Before coming to the computation of differential Galois groups, we first consider the simpler problem of factoring L in section 3. We recall that there exists a non-trivial factorization of L if and only if the Galois group of L admits a non-trivial invariant subspace. By using computations with limited precision, we show how to use this criterion in order to compute candidate factorizations or a proof that there exist no factorizations. It is easy to check *a posteriori* whether a candidate factorization is correct, so we obtain a factorization algorithm by increasing the precision until we obtain a correct candidate or a proof that there are no factorizations.

In section 4 we consider the problem of computing the differential Galois group of L. Using the results from section 2, it suffices to show how to compute the algebraic closure of a matrix group $\mathcal{G}$ generated by a finite number of given elements. A theoretical solution for this problem based on Gröbner basis techniques has been given in [DJK03]. The main idea behind the present algorithm is similar, but more emphasis is put on efficiency (in contrast to generality).

First of all, in our context of complex numbers with arbitrary precisions, we may use the LLL-algorithm for the computation of linear and multiplicative dependencies [LLL82]. Secondly, the connected component of $\mathcal{G}$ is represented as the exponential of a Lie algebra $\mathcal{L}$ given by a basis. Computations with such Lie algebras essentially boil down to linear algebra. Finally, we use classical techniques for finite groups in order to represent and compute with the elements in $\mathcal{G}/\mathrm{e}^{\mathcal{L}}$ [Sim70, Sim71, MO95]. Moreover, we will present an algorithm for non-commutative lattice reduction, similar to the LLL-algorithm, for the efficient computation with elements in $\mathcal{G}/\mathrm{e}^{\mathcal{L}}$ near the identity.

The algorithms in section 4 are all done using a fixed precision. Although we do prove that we really compute the Galois group when using a sufficiently large precision, it is not clear *a priori* how to find such a "sufficient precision". Nevertheless, we have already seen in section 3 that it is often possible to check the correctness of the result *a posteriori*, especially when we are not interested in the Galois group $\mathcal{G}$ itself, but only in some information provided by $\mathcal{G}$. Also, it might be possible to reduce the amount of dependence on "transcendental arguments" in the algorithm modulo a further development of our ideas. Some hints are given in the last section.

REMARK 1. The author first suggested the main approach behind this paper during his visit at the MSRI in 1998. The outline of the algorithm in section 4.5 came up in a discussion with Harm Derksen (see also [DJK03]). The little interest manifested by specialists in effective differential Galois theory for this approach is probably due to the fact that current computer algebra systems have very poor support for analytic computations. We hope that the present article will convince people to put more effort in the implementation of such algorithms. We started such an effort [vdHea05], but any help would be appreciated. Currently, none of the algorithms presented in this paper has been implemented.

2. PRELIMINARIES

2.1. Notations

Throughout this paper, we will use the following notations:

$\mathbb{K} \subseteq \mathbb{C}$. An algebraically closed field of constants.

$\mathsf{Mat}_n(\mathbb{K})$. The algebra of $n \times n$ matrices with coefficients in $\mathbb{K}$.

$\mathsf{GL}_n(\mathbb{K})$. The subgroup of $\mathsf{Mat}_n(\mathbb{K})$ of invertible matrices.

$\mathsf{Vect}(S)$. The vector space generated by a subset S of a larger vector space.

$\mathsf{Alg}(S)$. The algebra generated by a subset S of a larger algebra.

Vectors are typeset in bold face $\boldsymbol{v} = (v_1, \ldots, v_n)$ and we use the following vector notations:

$$\begin{aligned} \boldsymbol{v} \cdot \boldsymbol{w} &= v_1 w_1 + \cdots + v_n w_n \\ \boldsymbol{v}^{\boldsymbol{k}} &= v_1^{k_1} \cdots v_n^{k_n} \end{aligned}$$

Matrices $M \in \mathsf{Mat}_n(\mathbb{K})$ will also be used as mappings $M: \mathbb{K}^n \to \mathbb{K}^n; \boldsymbol{v} \mapsto M \boldsymbol{v}$. When making a base change in $\mathbb{K}^n$, we understand that we perform the corresponding transformations $M \to P M P^{-1}$ on all matrices under consideration. We denote $\mathrm{Diag}(X_1, \ldots, X_p)$ for the diagonal matrix with entries $X_1, \ldots, X_p$. The X_i may either be scalars or square matrices. Given a matrix $M \in \mathsf{Mat}_n(\mathbb{K})$ and a vector $\boldsymbol{v} \in \mathbb{K}^n$, we write $\boldsymbol{v}^M$ for the vector $\boldsymbol{w}$ with $w_i = v_1^{M_{i,1}} \cdots v_n^{M_{i,n}}$ for all i.

2.2. Differential Galois groups

Consider a monic linear differential operator $L = \partial^n + L_{n-1} \partial + \cdots + L_0 \in \mathcal{F}[\partial]$, where $\mathbb{K}$ is an algebraically closed subfield of $\mathbb{C}$ and $\mathcal{F} = \mathbb{K}(z)$. We will denote by $\mathcal{S} = \mathcal{S}_L \subseteq \mathbb{K} \cup \{\infty\}$ the finite set of singularities of L (in the case of ∞, one considers the transformation $z \mapsto z^{-1}$). A *Picard-Vessiot extension* of $\mathcal{F}$ is a differential field $\mathcal{K} \supseteq \mathcal{F}$ such that

PV1. $\mathcal{K} = \mathcal{F}\langle h_1, \ldots, h_n \rangle$ is differentially generated by $\mathcal{F}$ and a basis of solutions $\boldsymbol{h} = (h_1, \ldots, h_n) \in \mathcal{K}^n$ to the equation $L f = 0$.

PV2. $\mathcal{K}$ has $\mathbb{K}$ as its field of constants.

A Picard-Vessiot extension always exists: given a point $z_0 \in \mathbb{K} \setminus \mathcal{S}$ and $i \in \{1, \ldots, n\}$, let h_i be the unique solution to $Lf = 0$ with $h_i^{(j)}(z_0) = \delta_{i,j+1}$ for $j \in \{0, \ldots, n-1\}$. We call $\boldsymbol{h} = \boldsymbol{h}^{z_0} = (h_1, \ldots, h_n)$ the *canonical basis* for the solution space of $Lf = 0$ at z_0, and regard $\boldsymbol{h}$ as a column vector. Taking $\mathcal{K} = \mathcal{F}\langle h_1, \ldots, h_n\rangle$, the condition **PV2** is trivially satisfied since $\boldsymbol{h}(z_0 + \varepsilon) \in \mathbb{K}[[\varepsilon]] \subseteq \mathbb{K}((\varepsilon))$ and the constant field of $\mathbb{K}((z))$ is $\mathbb{K}$.

Let $\mathcal{K}$ be a Picard-Vessiot extension of $\mathcal{F}$ and let $\boldsymbol{h} \in \mathcal{K}^n$ be as in **PV1**. The *differential Galois group* $\mathcal{G}_{\mathcal{K}/\mathcal{F}}$ of the extension $\mathcal{K}/\mathcal{F}$ is the group of differential automorphisms which leave $\mathcal{F}$ pointwise invariant. It is classical [Kol73] that $\mathcal{G}_{\mathcal{K}/\mathcal{F}}$ is independent (up to isomorphism) of the particular choice of the Picard-Vessiot extension $\mathcal{K}$.

Given an automorphism $\sigma \in \mathcal{G}_{\mathcal{K}/\mathcal{F}}$, any solution f to $Lf = 0$ is sent to another solution. In particular, there exists a unique matrix $M = M_{\sigma,\boldsymbol{h}} \in \mathsf{GL}_n(\mathbb{K})$ with $\sigma h_i = M h_i := \sum_{j=1}^n M_{i,j} h_j$ for all i. This yields an embedding $\rho_{\boldsymbol{h}}$ of $\mathcal{G}_{\mathcal{K}/\mathcal{F}}$ into $\mathsf{GL}_n(\mathbb{K})$ and we define $\mathcal{G}_{L,\boldsymbol{h}} := \rho_{\boldsymbol{h}}(\mathcal{G}_{\mathcal{K}/\mathcal{F}})$. Conversely, $M \in \mathsf{GL}_n(\mathbb{K})$ belongs to $\mathcal{G}_{L,\boldsymbol{h}}$ if every differential relation $P(h_1, \ldots, h_n) = 0$ satisfied by $h_1, \ldots, h_n$ is also satisfied by $M h_1, \ldots, M h_n$ (with $P \in \mathbb{K}\{F_1, \ldots, F_n\}$). Since this assumption constitutes an infinite number of algebraic conditions on the coefficients of M, it follows that $\mathcal{G}_{L,\boldsymbol{h}}$ is a Zariski closed algebraic matrix group. Whenever $\boldsymbol{g} = P\boldsymbol{h}$ is another basis, we obtain the same matrix group $\mathcal{G}_{L,\boldsymbol{g}} = P\mathcal{G}_{L,\boldsymbol{h}} P^{-1}$ up to conjugation.

Assume now that $\hat{\mathbb{K}} \supseteq \mathbb{K}$ is a larger algebraically closed subfield of $\mathbb{C}$. Then the field $\hat{\mathcal{K}} = \hat{\mathbb{K}}(z)\langle h_1, \ldots, h_n\rangle = \mathcal{K} \otimes \hat{\mathbb{K}}$ is again a Picard-Vessiot extension of $\hat{\mathbb{K}}(z)$. Furthermore, the Ritt-Raudenbush theorem [Rit50] implies that the perfect differential ideal of all $P \in \mathbb{K}\{F_1, \ldots, F_n\}$ with $P(h_1, \ldots, h_n) = 0$ is finitely generated, say by $G_1, \ldots, G_k$. But then $G_1, \ldots, G_k$ is still a finite system of generators of the perfect differential ideal of all $P \in \hat{\mathbb{K}}\{F_1, \ldots, F_n\}$ with $P(h_1, \ldots, h_n) = 0$. Consequently, $\hat{\mathcal{G}}_{L,\boldsymbol{g}} \subseteq \mathsf{GL}_n(\hat{\mathbb{K}})$ (i.e. as an algebraic group over $\hat{\mathbb{K}}$) is determined by the same algebraic equations as $\mathcal{G}_{L,\boldsymbol{g}}$. We conclude that $\mathcal{G}_{L,\boldsymbol{h}} = \hat{\mathcal{G}}_{L,\boldsymbol{h}} \cap \mathsf{GL}_n(\mathbb{K})$.

Let $\mathcal{K}$ be a Picard-Vessiot extension of $\mathcal{F}$. Any differential field $\mathcal{L}$ with $\mathcal{F} \subseteq \mathcal{L} \subseteq \mathcal{K}$ naturally induces an algebraic subgroup $\mathcal{L}' \subseteq \mathcal{G}_{\mathcal{K}/\mathcal{F}}$ of automorphisms of $\mathcal{K}$ which leave $\mathcal{L}$ fixed. Inversely, any algebraic subgroup $\mathcal{H}$ of $\mathcal{G}_{\mathcal{K}/\mathcal{L}}$ gives rise to the differential field $\mathcal{H}'$ with $\mathcal{F} \subseteq \mathcal{H}' \subseteq \mathcal{K}$ of all elements which are invariant under the action of $\mathcal{H}$. We say that $\mathcal{L}$ (resp. $\mathcal{H}$) is *closed* if $\mathcal{L} = \mathcal{L}''$ (resp. $\mathcal{H}'' = \mathcal{H}$). In that case, the extension $\mathcal{L}/\mathcal{F}$ is said to be *normal*, i.e. every element in $\mathcal{L} \setminus \mathcal{F}$ is moved by an automorphism of $\mathcal{L}$ over $\mathcal{F}$. The main theorem from differential Galois theory states that the Galois correspondences are bijective [Kap57, Theorem 5.9].

THEOREM 2. *With the above notations:*

a) *The correspondences $\mathcal{L} \mapsto \mathcal{L}'$ and $\mathcal{H} \mapsto \mathcal{H}'$ are bijective.*

b) *The group $\mathcal{H}$ is a closed normal subgroup of $\mathcal{G}_{\mathcal{K}/\mathcal{F}}$ if and only if the extension $\mathcal{H}'/\mathcal{F}$ is normal. In that case, $\mathcal{G}_{\mathcal{K}/\mathcal{F}}/\mathcal{H} \cong \mathcal{G}_{\mathcal{H}'/\mathcal{F}}$.*

COROLLARY 3. *Let $f \in \mathcal{F}\langle h_1, \ldots, h_n\rangle$. If $Mf = f$ for all $M \in \mathcal{G}_{L,\boldsymbol{h}}$, then $f \in \mathcal{F}$.*

2.3. Monodromy

Consider a continuous path γ on $\mathbb{C} \cup \{\infty\} \setminus \mathcal{S}$ from $z_0 \in \mathbb{K}$ to $z_1 \in \mathbb{K}$. Then analytic continuation of the canonical basis $\boldsymbol{h}^{z_0}$ at z_0 along γ yields a basis of solutions to $Lf = 0$ at z_1. The matrix $\Delta_\gamma \in \mathsf{GL}_n(\mathbb{K})$ with

$$\boldsymbol{h}^{z_1} = \Delta_\gamma \boldsymbol{h}^{z_0} \tag{1}$$

is called the *connection matrix* or *transition matrix* along γ. In particular, if $z_1 = z_0$, then we call Δ_γ a *monodromy matrix* based in z_0. We clearly have

$$\Delta_{\gamma_2 \circ \gamma_1} = \Delta_{\gamma_2} \Delta_{\gamma_1}$$

for the composition of paths, so the monodromy matrices based in z_0 form a group Mono_{z_0} which is called the *monodromy group*. Given a path γ from z_0 to z_1, we notice that $\mathsf{Mono}_{z_1} = \Delta_\gamma \, \mathsf{Mono}_{z_0} \, \Delta_\gamma^{-1}$. Since any differential relation satisfied by $\boldsymbol{h}^{z_0}$ is again satisfied by its analytic continuation along γ, we have $\mathsf{Mono}_{z_0} \subseteq \mathcal{G}_{L,\boldsymbol{h}^{z_0}}$ and $\mathcal{G}_{L,\boldsymbol{h}^{z_1}} = \Delta_\gamma \mathcal{G}_{L,\boldsymbol{h}^{z_0}} \Delta_\gamma^{-1}$.

REMARK 4. The definition of transition matrices can be slightly changed depending on the purpose [vdH05b, Section 4.3.1]: when interpreting $\boldsymbol{h}^{z_0}$ and $\boldsymbol{h}^{z_1}$ as row vectors, then (1) has to be transposed. The roles of $\boldsymbol{h}^{z_0}$ and $\boldsymbol{h}^{z_1}$ may also be interchanged modulo inversion of Δ_γ.

Now assume that L admits a singularity at 0 (if $\mathcal{S} \neq \varnothing$ then we may reduce to this case modulo a translation; singularities at infinity may be brought back to zero using the transformation $z \to z^{-1}$). It is well-known [Fab85, vH96] that $L f$ admits a computable formal basis of solutions of the form

$$f = (f_0(\sqrt[p]{z}) + \cdots + f_{n-1}(\sqrt[p]{z}) \log^{n-1} z) \, z^\alpha \, \mathrm{e}^{P(\sqrt[p]{z})}, \tag{2}$$

with $h_0, \ldots, h_{n-1} \in \mathbb{K}[[z]]$, $p \in \mathbb{N}^>$, $\alpha \in \mathbb{K}$ and $P \in \mathbb{K}[z]$. We will denote by $\mathbb{S}$ the set of finite sums of expressions of the form (2). We may see $\mathbb{S}$ as a differential subring of a formal differential field of "complex transseries" $\mathbb{T}$ [vdH01a] with constant field $\mathbb{C}$.

We recall that transseries in $\mathbb{T}$ are infinite linear combinations $f = \sum_{\mathfrak{m} \in \mathfrak{T}} f_\mathfrak{m} \mathfrak{m}$ of "transmonomials" with "grid-based support". The set $\mathfrak{T}$ of transmonomials forms a totally ordered vector space for exponentiation by reals and the asymptotic ordering $\preccurlyeq$. In particular, each non-zero transseries f admits a unique dominant monomial $\mathfrak{d}_f$. It can be shown [vdH01a] that there exists a unique basis $\boldsymbol{h} = (h_1, \ldots, h_n)$ of solutions to $L f = 0$ of the form (2), with $h_1 \prec \cdots \prec h_n$ and $(h_i)_{\mathfrak{d}(h_j)} = \delta_{i,j}$ for all $i, j \in \{1, \ldots, n\}$. We call $\boldsymbol{h}^0 = \boldsymbol{h}$ the canonical basis of solutions in 0 and there is an algorithm which computes $\boldsymbol{h}$ as a function of L.

Let $\mathbb{L}$ be the subset of $\mathbb{S}$ of all finite sums of expressions of the form (2) with $P = 0$. Then any $f \in \mathbb{S}$ can uniquely be written as a finite sum $f = \sum_{\mathfrak{e} \in \mathfrak{E}} f_\mathfrak{e} \mathfrak{e}$, where $\mathfrak{E} = \exp(\bigcup_p \mathbb{C}[\sqrt[p]{z}])$. Let Expo_0 be the group of all automorphisms $\sigma: \mathbb{S} \to \mathbb{S}$ for which there exists a mapping $\lambda: \mathfrak{E} \to \mathbb{K}^{\neq}; \mathfrak{e} \mapsto \lambda_\mathfrak{e}$ with $\sigma(f) = \sum_{\mathfrak{e} \in \mathfrak{E}} \lambda_\mathfrak{e} f_\mathfrak{e} \mathfrak{e}$ for all $f \in \mathbb{S}$. Then every $\sigma \in \mathbb{S}$ preserves differentiation and maps the Picard-Vessiot extension $\mathcal{K} = \mathcal{F}\langle h_1, \ldots, h_n \rangle$ of $\mathcal{F}$ into itself. In particular, the restriction $\mathsf{Expo}_{0,\boldsymbol{h}}$ of Expo_0 to $\mathcal{K}$ is a subset of $\mathcal{G}_{L,\boldsymbol{h}}$.

PROPOSITION 5. *Assume that $f \in \mathbb{S}$ is fixed under Expo_0. Then $f \in \mathbb{L}$.*

PROOF. Assume that $f \notin \mathbb{L}$ and let $\mathfrak{e} \in \mathfrak{E}$ be an exponential with $f_\mathfrak{e} \neq 0$. Let $\mathcal{H}$ be a supplement of the $\mathbb{Q}$-vector space $\log \mathfrak{E}$. Let $\sigma: \mathbb{S} \to \mathbb{S}$ be the mapping in Expo_0 which sends $\mathfrak{e}^\alpha \mathfrak{f}$ to $\mathrm{e}^\alpha \mathfrak{e}^\alpha \mathfrak{f}$ for each $\alpha \in \mathbb{Q}$ and $\mathfrak{f} \in \exp \mathcal{H}$. Then we clearly have $\sigma(f) \neq f$. □

Let $\mathfrak{e}_1, \ldots, \mathfrak{e}_n$ be the set of exponents corresponding to the exponents of the elements of the canonical basis $\boldsymbol{h}^0$. Using linear algebra, we may compute a multiplicatively independent set $\mathfrak{f}_1, \ldots, \mathfrak{f}_r \in \mathfrak{e}_1^\mathbb{Q} \cdots \mathfrak{e}_n^\mathbb{Q}$ such that $\mathfrak{e}_i = \mathfrak{f}_1^{\beta_{i,1}} \cdots \mathfrak{f}_r^{\beta_{i,r}}$ for certain $\beta_{i,j} \in \mathbb{Z}$ and all i.

PROPOSITION 6. *With the above notations, the algebraic group $\mathsf{Expo}_{0,\boldsymbol{h}}$ is generated by the matrices $\mathrm{Diag}(\lambda^{\beta_{i,1}}, \ldots, \lambda^{\beta_{i,n}})$ where $\lambda \in \mathbb{K}^{\neq} \setminus \{\mu: \exists n, \mu^n = 1\}$ is chosen arbitrarily.*

PROOF. Let $\mathcal{E}$ be the group generated by the matrices $\mathrm{Diag}(\lambda^{\beta_{i,1}}, \ldots, \lambda^{\beta_{i,n}})$. Notice that each individual matrix $\mathrm{Diag}(\lambda^{\beta_{i,1}}, \ldots, \lambda^{\beta_{i,n}})$ generates $\mathcal{S} = \{\mathrm{Diag}(\alpha^{\beta_{i,1}}, \ldots, \alpha^{\beta_{i,n}}) : \alpha \in \mathbb{K}^{\neq}\}$: assuming $\mathfrak{e}_i \neq 1$, the variety $\mathcal{S}$ is irreducible of dimension 1 and $\mathrm{Diag}(\lambda^{\beta_{i,1}}, \ldots, \lambda^{\beta_{i,n}})$ is not contained in an algebraic group of dimension 0. Now any $\sigma \in \mathsf{Expo}_{0,\boldsymbol{h}}$ is a diagonal matrix $\mathrm{Diag}(\lambda_{\mathfrak{e}_1}, \ldots, \lambda_{\mathfrak{e}_n})$ for some multiplicative mapping $\lambda: \mathfrak{E} \mapsto \mathbb{K}^{\neq}$. Hence

$$\mathrm{Diag}(\lambda_{\mathfrak{e}_1}, \ldots, \lambda_{\mathfrak{e}_n}) = \mathrm{Diag}(\lambda_{\mathfrak{f}_1}^{\beta_{1,1}}, \ldots, \lambda_{\mathfrak{f}_1}^{\beta_{n,1}}) \cdots \mathrm{Diag}(\lambda_{\mathfrak{f}_r}^{\beta_{1,r}}, \ldots, \lambda_{\mathfrak{f}_r}^{\beta_{n,r}}) \in \mathcal{E}.$$

Conversely, each element

$$\sigma \in \mathrm{Diag}(\alpha_1^{\beta_{1,1}}, \ldots, \alpha_1^{\beta_{n,1}}) \cdots \mathrm{Diag}(\alpha_r^{\beta_{1,r}}, \ldots, \alpha_r^{\beta_{n,r}}) \in \mathcal{E}$$

determines a multiplicative mapping $\lambda: \mathfrak{f}_1^{\mathbb{Z}} \cdots \mathfrak{f}_r^{\mathbb{Z}} \to \mathbb{K}^{\neq}; f_1^{k_1} \cdots f_r^{k_r} \mapsto \alpha_1^{k_1} \cdots \alpha_r^{k_r}$ which may be further extended to $\mathfrak{E}$ using Zorn's lemma and the fact that $\mathbb{K}$ is algebraically closed. It follows that $\sigma \in \mathsf{Expo}_{0,\boldsymbol{h}}$. □

Assume that $2\,\pi\,\mathrm{i} \in \mathbb{K}$ and let $M_0: \mathbb{S} \to \mathbb{S}$ be the transformation which sends $\log z$ to $\log z + 2\,\pi\,\mathrm{i}$, z^α to $\mathrm{e}^{2\pi\mathrm{i}\alpha}\, z^\alpha$ and $\mathrm{e}^{P(\sqrt[p]{z})}$ to $\mathrm{e}^{M_0(P(\sqrt[p]{z}))}$. Then σ preserves differentiation, so any solution to $L f = 0$ of the form (2) is sent to another solution of the same form. In particular, there exists a matrix $\Delta_{\circlearrowleft 0}$ with $M_0\, \boldsymbol{h} = \Delta_{\circlearrowleft 0}\, \boldsymbol{h}$, called the *formal monodromy matrix* around 0. We have $\Delta_{\circlearrowleft 0} \in \mathcal{G}_{L,\boldsymbol{h}}$.

PROPOSITION 7. *Assume that $f \in \mathbb{S}$ is fixed under* Expo_0 *and M_0. Then $f \in \mathbb{K}((z))$.*

PROOF. We already know that $f \in \mathbb{L}$. Interpreting $f = c_k \log^k z + \cdots + c_0$ as a polynomial in $\log z$ with $k > 0 \Rightarrow c_k \neq 0$, we must have $k = 0$ since

$$M_0(f) - f = 2\,\pi\,\mathrm{i}\, k\, c_k \log^{k-1} z + \cdots = 0.$$

Consequently, f is of the form $f = \sum_{\alpha \in \mathbb{K}} f_\alpha\, z^\alpha$ and

$$M_0(f) = \sum_{\alpha \in \mathbb{K}} \mathrm{e}^{2\pi\mathrm{i}\alpha}\, f_\alpha\, z^\alpha = \sum_{\alpha \in \mathbb{K}} f_\alpha\, z^\alpha = f(z),$$

We conclude that $\mathrm{e}^{2\pi\mathrm{i}\alpha} = 1$ for every $\alpha \in \mathbb{K}$ with $f_\alpha \neq 0$, whence $f \in \mathbb{C}((z))$. □

2.4. The process of accelero-summation

Let $\mathbb{C}[[z^{\mathbb{Q}^>}]]$ be the differential $\mathbb{C}$-algebra of infinitesimal Puiseux series in z for $\delta = z\,\partial$ and consider a formal power series solution $\tilde{f} \in \mathbb{O} = \mathbb{C}[[z^{\mathbb{Q}^>}]][\log z]$ to $L\,\tilde{f} = 0$. The process of accelero-summation enables to associate an analytic meaning f to $\tilde{f}$ in a sector near the origin of the Riemann surface $\dot{\mathbb{C}}$ of log, even in the case when $\tilde{f}$ is divergent. Schematically speaking, we obtain f through a succession of transformations:

$$\begin{tikzcd} \tilde{f} \arrow[d, "\tilde{\mathcal{B}}_{z_1}"'] & & & & & f \\ \hat{f}_1 \arrow[r, "\mathcal{A}^{\alpha_1}_{z_1 \to z_2}"'] & \hat{f}_2 \arrow[r] & \cdots \arrow[r] & \hat{f}_{p-1} \arrow[rr, "\mathcal{A}^{\alpha_{p-1}}_{z_{p-1} \to z_p}"'] & & \hat{f}_n \arrow[u, "\mathcal{L}^{\alpha_p}_{z_p}"'] \end{tikzcd} \tag{3}$$

Each $\hat{f}_i$ is a "resurgent function" which realizes $\tilde{f}_i(z_i) = \tilde{f}(z)$ in the "convolution model" with respect to the i-th "critical time" $z_i = \sqrt[k_i]{z}$ (with $k_i \in \mathbb{Q}^>$ and $k_1 > \cdots > k_p$). In our case, $\hat{f}_i$ is an analytic function which admits only a finite number of singularities above $\mathbb{C}$. In general, the singularities of a resurgent function are usually located on a finitely generated grid. Let us describe the transformations $\tilde{\mathcal{B}}$, $\mathcal{A}^{\alpha_i}_{z_i \to z_{i+1}}$ and $\mathcal{L}^{\alpha_p}_{z_p}$ in more detail.

THE BOREL TRANSFORM. We start by applying the *formal Borel transform* to the series $\tilde{f}_1(z_1) = \tilde{f}(z) = \sum_{\sigma,r} \tilde{f}_{1,\sigma,r}\, z_1^\sigma \log^r z_1 \in \mathbb{C}[[z_1^{\mathbb{Q}^>}]][\log z_1]$. This transformation sends each $z_1^\sigma \log^r z_1$ to

$$(\tilde{\mathcal{B}}_{z_1} z_1^\sigma \log^r z_1)(\zeta_1) = \zeta_1^{\sigma-1} \sum_{i=0}^{r} \binom{k}{i} \gamma^{(r-i)}(\sigma) \log^i \zeta_1,$$

where $\gamma(\sigma) = 1/\Gamma(\sigma)$, and extends by strong linearity:

$$\hat{f}_1(\zeta_1) = (\tilde{\mathcal{B}}_{z_1} \tilde{f}_1)(\zeta_1) = \sum_{\substack{\sigma \in \mathbb{Q}^> \\ r \in \mathbb{N}}} \tilde{f}_{1,r,\sigma} (\tilde{\mathcal{B}}_{z_1} z_1^\sigma \log^r z_1)(\zeta_1),$$

The result is a formal series $\hat{f}_1 \in \zeta_1^{-1}\, \mathbb{C}[[\zeta_1^{\mathbb{Q}^>}]][\log \zeta_1]$ in ζ_1 which converges near the origin of $\dot{\mathbb{C}}$. The formal Borel transform is a morphism of differential algebras which sends multiplication to the convolution product, i.e. $\tilde{\mathcal{B}}_{z_1}(f g) = (\tilde{\mathcal{B}}_{z_1} f) * (\tilde{\mathcal{B}}_{z_1} g)$.

ACCELERATIONS. Given $i < p$, the function $\hat{f}_i$ is defined near the origin of $\dot{\mathbb{C}}$, can be analytically continued on the axis $\mathrm{e}^{\alpha_i \mathrm{i}}\, \mathbb{R}^> \subseteq \dot{\mathbb{C}}$, and admits a growth of the form $\hat{f}_i(\zeta_i) = \exp O(|\zeta_i|^{k_i/(k_i - k_{i+1})})$ at infinity. The next function $\hat{f}_{i+1}$ is obtained from $\hat{f}_i$ by an *acceleration* of the form

$$\hat{f}_{i+1}(\zeta_{i+1}) = (\mathcal{A}_{z_i \to z_{i+1}}^{\alpha_i} \hat{f}_i)(\zeta_{i+1}) = \int_{\zeta_i \in \mathrm{e}^{\alpha_i \mathrm{i}} \mathbb{R}^>} K_{k_i, k_{i+1}}(\zeta_i, \zeta_{i+1})\, \hat{f}_i(\zeta_i)\, \mathrm{d}\zeta_i,$$

where the acceleration kernel $K_{k_i, k_{i+1}}$ is given by

$$\begin{aligned} K_{k_i,k_{i+1}}(\zeta_i, \zeta_{i+1}) &= \frac{1}{\zeta_{i+1}} K_{k_{i+1}/k_i}\left(\frac{\zeta_i}{\zeta_{i+1}^{k_{i+1}/k_i}}\right) \\ K_\lambda(\zeta) &= \frac{1}{2\pi \mathrm{i}} \int_{c-\infty \mathrm{i}}^{c+\infty \mathrm{i}} \mathrm{e}^{z - \zeta z^\lambda} \mathrm{d}z. \end{aligned} \tag{4}$$

For large ζ on an axis with $|\arg \zeta| < (1 - \lambda)\, \pi/2$, it can be shown that $K_\lambda(\zeta) \leqslant \exp(-C\, |\zeta|^{1/(1-\lambda)})$ for some constant $C > 0$. Assuming that α_{i+1} satisfies

$$|k_{i+1}\, \alpha_{i+1} - k_i\, \alpha_i| < (k_i - k_{i+1})\, \pi/2, \tag{5}$$

it follows that the acceleration $\hat{f}_{i+1}$ of $\hat{f}_i$ is well-defined for small ζ_{i+1} on $\mathrm{e}^{\alpha_{i+1}}\, \mathbb{R}^>$. The set $\mathcal{D}_i \subseteq \mathbb{R}$ of directions α such $\hat{f}_i$ admits a singularity on $\mathrm{e}^{\alpha \mathrm{i}}\, \mathbb{R}^>$ is called the set of *Stokes directions*. Accelerations are morphisms of differential $\mathbb{C}$-algebras which preserve the convolution product.

THE LAPLACE TRANSFORM. The last function $\hat{f}_p$ is defined near the origin of $\dot{\mathbb{C}}$, can be analytically continued on the axis $\mathrm{e}^{\alpha_i \mathrm{i}}\, \mathbb{R}^> \subseteq \dot{\mathbb{C}}$ and admits at most exponential growth at infinity. The function f is now obtained using the analytic *Laplace transform*

$$f(z) = f_p(z_p) = (\mathcal{L}_{z_p}^{\alpha_p} \hat{f}_p)(z_p) = \int_{\zeta_p \in \mathrm{e}^{\alpha_p \mathrm{i}} \mathbb{R}^>} \hat{f}_p(\zeta_p)\, \mathrm{e}^{-\zeta_p / z_p}\, \mathrm{d}\zeta_p.$$

On an axis with

$$|\arg z_p - \alpha_p| < \pi/2, \tag{6}$$

the function f_p is defined for all sufficiently small z_p. The set $\mathcal{D}_p$ of Stokes directions is defined in a similar way as in the case of accelerations. The Laplace transform is a morphism of differential $\mathbb{C}$-algebras which is inverse to the Borel transform and sends the convolution product to multiplication.

REMARK 8. Intuitively speaking, one has $\mathcal{A}_{z_i \to z_{i+1}}^{\alpha_i} = \mathcal{B}_{z_{i+1}} \circ \mathcal{L}_{z_i}^{\alpha_i}$.

Given critical times $k_1 > \cdots > k_p$ in $\mathbb{Q}^>$ and directions $\alpha_1, \ldots, \alpha_p$ satisfying (5), we say that a formal power series $\tilde{f} \in \tilde{\mathbb{O}}$ is *accelero-summable* in the multi-direction $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_p)$ if the above scheme yields an analytic function $f(z)$ near the origin of any axis on $\dot{\mathbb{C}}$ satisfying (6). We denote the set of such power series by $\mathbb{O}_{\boldsymbol{k},\boldsymbol{\alpha}}$, where $\boldsymbol{k} = (k_1, \ldots, k_p)$. Inversely, given $\tilde{f} \in \mathbb{O}$, we denote by $\operatorname{dom}_{\mathrm{as}} \tilde{f}$ the set of all triples $\gamma = (\boldsymbol{k}, \boldsymbol{\alpha}, z)$ such that $\tilde{f} \in \mathbb{O}_{\boldsymbol{k},\boldsymbol{\alpha}}$ and so that $f(z)$ is well-defined. In that case, we write $f = \operatorname{sum}_{\boldsymbol{k},\boldsymbol{\alpha}}$ and $f(z) = \tilde{f}(\gamma)$.

The set $\mathbb{O}_{\boldsymbol{k},\boldsymbol{\alpha}}$ forms a differential subring of $\mathbb{O}$ and the map $\tilde{f} \mapsto f$ for $\tilde{f} \in \mathbb{O}_{\boldsymbol{k},\boldsymbol{\alpha}}$ is injective. If $\boldsymbol{k}'$ and $\boldsymbol{\alpha}'$ are obtained from $\boldsymbol{k}$ and $\boldsymbol{\alpha}$ by inserting a new critical time and an arbitrary direction, then we have $\mathbb{O}_{\boldsymbol{k},\boldsymbol{\alpha}} \subsetneq \mathbb{O}_{\boldsymbol{k}',\boldsymbol{\alpha}'}$. In particular, $\mathbb{O}_{\boldsymbol{k},\boldsymbol{\alpha}}$ contains $\mathbb{O}_{\mathrm{cv}} = \mathbb{C}\{\{z^{\mathbb{Q}^>}\}\}[\log z]$, where $\mathbb{C}\{\{z^{\mathbb{Q}^>}\}\}$ denotes the ring of convergent infinitesimal Puiseux series. Let $\mathcal{R}_1 = \mathbb{R} \setminus \mathcal{D}_1 \subseteq \mathbb{R}, \ldots, \mathcal{R}_p = \mathbb{R} \setminus \mathcal{D}_p \subseteq \mathbb{R}$ be sets of directions such that each $\mathcal{D}_i$ is finite modulo 2π. Let $\boldsymbol{\mathcal{R}}$ be the subset of $\mathcal{R}_1 \times \cdots \times \mathcal{R}_p$ of multi-directions $\boldsymbol{\alpha}$ which verify (5). We denote $\mathbb{O}_{\boldsymbol{k},\boldsymbol{\mathcal{R}}} = \bigcap_{\boldsymbol{\alpha} \in \boldsymbol{\mathcal{R}}} \mathbb{O}_{\boldsymbol{k},\boldsymbol{\alpha}}$, $\mathbb{O}_{\boldsymbol{k}} = \bigcup_{\boldsymbol{\mathcal{R}}} \mathbb{O}_{\boldsymbol{k},\boldsymbol{\mathcal{R}}}$ and $\mathbb{O}_{\mathrm{as}} = \bigcup_{\boldsymbol{k}} \mathbb{O}_{\boldsymbol{k}}$.

Taking $\mathbb{K} = \mathbb{C}$, the notion of accelero-summation extends to formal expressions of the form (2) and more general elements of $\mathbb{S}$ as follows. Given $\tilde{g} \in \mathbb{O}_{\boldsymbol{k},\boldsymbol{\alpha}}$, $\sigma \in \mathbb{C}$, $\mathfrak{e} = \mathrm{e}^{P(\sqrt[p]{z})} \in \mathfrak{E}$ and $\gamma = (\boldsymbol{k}, \boldsymbol{\alpha}, z) \in \operatorname{dom}_{\mathrm{as}} \tilde{g}$, we simply define $(\tilde{g}\, z^\sigma\, \mathfrak{e})(\gamma) = \tilde{g}(\gamma)\, z^\sigma\, \mathrm{e}^{P(\sqrt[p]{z})}$. It can be checked that this definition is coherent when replacing $\tilde{g}\, z^\sigma$ by $(z^k\, \tilde{g})\, z^{\sigma-k}$ for some $k \in \mathbb{Q}$. By linearity, we thus obtain a natural differential subalgebra $\mathbb{S}_{\boldsymbol{k},\boldsymbol{\alpha}} \subseteq \mathbb{S}$ of accelero-summable transseries with critical times $\boldsymbol{k}$ and in the multi-direction $\boldsymbol{\alpha}$. We also have natural analogues $\mathbb{S}_{\boldsymbol{k}}$ and $\mathbb{S}_{\mathrm{as}}$ of $\mathbb{O}_{\boldsymbol{k}}$ and $\mathbb{O}_{\mathrm{as}}$.

The main result we need from the theory of accelero-summation is the following theorem [É87, Bra91].

THEOREM 9. *Let $\tilde{f} \in \mathbb{O}$ be a formal solution to $L \tilde{f} = 0$. Then $\tilde{f} \in \mathbb{O}_{\mathrm{as}}$.*

COROLLARY 10. *Let $\boldsymbol{h}^0 \in \mathbb{S}^n$ be the canonical basis of formal solutions to $L \tilde{f} = 0$ at the origin. We have $\boldsymbol{h}^0 \in \mathbb{S}_{\mathrm{as}}^n$.*

PROOF. Holonomy is preserved under multiplication with elements of $z^{\mathbb{C}}\, \mathfrak{E}$. □

REMARK 11. We have aimed to keep our survey of the accelero-summation process as brief as possible. It is more elegant to develop this theory using resurgent functions and resurgence monomials [É85, CNP93].

2.5. The Stokes phenomenon

We say that $\tilde{f} \in \mathbb{S}_{\boldsymbol{k},\boldsymbol{\mathcal{R}}}$ is *stable under Stokes morphisms* is for all $\boldsymbol{\alpha}, \boldsymbol{\beta} \in \boldsymbol{\mathcal{R}}$, there exists a $\tilde{g} \in \mathbb{S}_{\boldsymbol{k},\boldsymbol{\mathcal{R}}}$ with $\operatorname{sum}_{\boldsymbol{k},\boldsymbol{\alpha}} \tilde{f} = \operatorname{sum}_{\boldsymbol{k},\boldsymbol{\beta}} \tilde{g}$, and if the same property is recursively satisfied by $\tilde{g}$. We denote by $\check{\mathbb{S}}_{\boldsymbol{k},\boldsymbol{\mathcal{R}}}$ the differential subring of $\mathbb{S}_{\boldsymbol{k},\boldsymbol{\mathcal{R}}}$ which is stable under Stokes morphisms. The mappings $\Sigma_{\boldsymbol{k},\boldsymbol{\alpha},\boldsymbol{\beta}}: \tilde{f} \mapsto \operatorname{sum}_{\boldsymbol{k},\boldsymbol{\beta}}^{-1} \operatorname{sum}_{\boldsymbol{k},\boldsymbol{\alpha}} \tilde{f}$ will be called *Stokes morphisms* and we denote by $\mathsf{Sto}_{0,\boldsymbol{k},\boldsymbol{\mathcal{R}}}$ the group of all such maps.

PROPOSITION 12. *Assume that $\tilde{f} \in \mathbb{S}_{\boldsymbol{k},\boldsymbol{\mathcal{R}}}$ is fixed under $\mathsf{Sto}_{0,\boldsymbol{k},\boldsymbol{\mathcal{R}}}$. Then $\tilde{f}$ is convergent.*

PROOF. Assume that one of the $\hat{f}_i$ admits a singularity at $\omega = \rho\, \mathrm{e}^{\theta \mathrm{i}} \neq 0$ and choose i maximal and ρ minimal. Modulo the removal of unnecessary critical times, we may assume without loss of generality that $i = p$. Let $\boldsymbol{\alpha}$ with $\alpha_p = \theta$ be a multi-direction satisfying (5), such that $\alpha_i \in \mathcal{R}_i$ for all $i < p$. Then

$$
\begin{aligned}
\boldsymbol{\alpha}^< &= (\alpha_1, \ldots, \alpha_{p-1}, \alpha_p - \varepsilon) \in \boldsymbol{\mathcal{R}} \\
\boldsymbol{\alpha}^> &= (\alpha_1, \ldots, \alpha_{p-1}, \alpha_p + \varepsilon) \in \boldsymbol{\mathcal{R}}
\end{aligned}
$$

for all sufficiently small $\varepsilon > 0$. Now $g_p = \mathcal{L}_{z_p}^{\theta+\varepsilon} \hat{f}_p - \mathcal{L}_{z_p}^{\theta-\varepsilon} \hat{f}_p$ is obtained by integration around ω along the axis $\mathrm{e}^{\theta \mathrm{i}} \mathbb{R}^{>}$. By classical properties of the Laplace integral [CNP93, Pré I.2], the function g_p cannot vanish, since $\hat{f}_i$ admits a singularity in ω (if the Laplace integrals corresponding to both directions $\theta \pm \varepsilon$ coincide, then the Laplace transform can be analytically continued to a larger sector, which is only possible if $\hat{f}_i$ is analytic in a sector which contains both directions $\theta \pm \varepsilon$). We conclude that $g(z) = g_p(z_p) = (\mathrm{sum}_{\boldsymbol{k},\boldsymbol{\alpha}^>} - \mathrm{sum}_{\boldsymbol{k},\boldsymbol{\alpha}^<})(\tilde{f}) \neq 0$, so f is not fixed under $\mathsf{Sto}_{0,\boldsymbol{k},\boldsymbol{\mathcal{R}}}$. □

REMARK 13. Let $\boldsymbol{\mathcal{D}}$ be a set of multi-directions $\boldsymbol{\alpha}$ satisfying (5), with $\alpha_i \in \mathcal{D}_i$ for exactly one i, and so that for all $j \neq i$, we have either $\alpha_j = k_i \alpha_i / k_j$ or $k_i \alpha_i / k_j \in \mathcal{D}_j$ and $\alpha_j = k_i (\alpha_i \pm \varepsilon)/k_j$ for some small $\varepsilon > 0$. For every $\boldsymbol{\alpha} \in \boldsymbol{\mathcal{D}}$, we have

$$\begin{aligned} \boldsymbol{\alpha}^< &= (\alpha_1, \ldots, \alpha_{i-1}, \alpha_i - \varepsilon, \alpha_{i+1}, \ldots, \alpha_p) \in \boldsymbol{\mathcal{R}} \\ \boldsymbol{\alpha}^> &= (\alpha_1, \ldots, \alpha_{i-1}, \alpha_i + \varepsilon, \alpha_{i+1}, \ldots, \alpha_p) \in \boldsymbol{\mathcal{R}}. \end{aligned}$$

By looking more carefully at the proof of proposition 12, we observe that it suffices to assume that $\tilde{f}$ is fixed under all Stokes morphisms of the form $\Sigma_{\boldsymbol{k},\boldsymbol{\alpha}^<,\boldsymbol{\alpha}^>}$, instead of all elements in $\mathsf{Sto}_{0,\boldsymbol{k},\boldsymbol{\mathcal{R}}}$.

We say that $\boldsymbol{\alpha}, \boldsymbol{\beta} \in \boldsymbol{\mathcal{D}}$ are equivalent, if $\alpha_i - \beta_i \in 2\pi \mathrm{i}\, q_i$ for all i, where q_i is the denominator of k_i. We notice that $\boldsymbol{\mathcal{D}}$ is finite modulo this equivalent relation. We denote by $\boldsymbol{\mathcal{D}}^{\mathrm{gen}}$ a subset of $\boldsymbol{\mathcal{D}}$ with one element in each equivalence class.

Let us now come back to our differential equation $L f = 0$. Given $\gamma = (\boldsymbol{k}, \boldsymbol{\alpha}, z) \in \mathrm{dom}_{\mathrm{as}}\, \boldsymbol{h}^0 := \mathrm{dom}_{\mathrm{as}}\, h_1^0 \cap \cdots \cap \mathrm{dom}_{\mathrm{as}}\, h_n^0$, the map $\mathrm{sum}_{\boldsymbol{k},\boldsymbol{\alpha}}$ induces an isomorphism between $\mathsf{Vect}(\boldsymbol{h}^0)$ and $\mathsf{Vect}(\boldsymbol{h}^z)$. We denote by $\Delta_\gamma \in \mathsf{GL}_n(\mathbb{C})$ the unique matrix with $\boldsymbol{h}^z = \Delta_\gamma \mathrm{sum}_{\boldsymbol{k},\boldsymbol{\alpha}} \boldsymbol{h}^0$. Given a second $\gamma' = (\boldsymbol{k}, \boldsymbol{\alpha}', z) \in \mathrm{dom}_{\mathrm{as}} \boldsymbol{h}^0$, the vector $\mathrm{sum}_{\boldsymbol{k},\boldsymbol{\alpha}'}^{-1} \mathrm{sum}_{\boldsymbol{k},\boldsymbol{\alpha}} \boldsymbol{h}^0$ is again in $\mathbb{S}_{\boldsymbol{k},\boldsymbol{\mathcal{R}}}^n$, whence $\boldsymbol{h}^0 \in \check{\mathbb{S}}_{\boldsymbol{k},\boldsymbol{\mathcal{R}}}$ by repeating the argument. In particular, the Stokes morphism $\Sigma_{\boldsymbol{k},\boldsymbol{\alpha},\boldsymbol{\alpha}'}$ induces the *Stokes matrix* $\Delta_{(0,\boldsymbol{k},\boldsymbol{\alpha}\to\boldsymbol{\alpha}')} = \Delta_{\gamma'}^{-1} \Delta_\gamma$.

We are now in the position that we can construct a finite set $\mathcal{M}$ of generators for the Galois group $\mathcal{G}_{L,\boldsymbol{h}^{z_0}}$ in a regular point $z_0 \in \mathbb{C} \cup \{+\infty\} \setminus \mathcal{S}$.

Algorithm `Compute_generators`(L, z_0)
Input: an operator $L \in \mathcal{F}[\partial]^{\neq}$ and a regular point $z_0 \in \mathbb{C} \cup \{+\infty\} \setminus \mathcal{S}$
Output: a set $\mathcal{M}$ of generators for $\mathcal{G}_{L,\boldsymbol{h}^{z_0}}$

$\mathcal{M} := \varnothing$
for each $z_i \in \mathcal{S}$ **do**

- Reduce to the case when $z_i = 0$ modulo a suitable transformation of the form $z \mapsto z + c$ or $z \mapsto z^{-1}$.
- Let γ_i be an arbitrary path $(\boldsymbol{k}, \boldsymbol{\alpha}, u_i) \in \mathrm{dom}_{\mathrm{as}}\, \boldsymbol{h}^{z_i}$ from z_i to a point u_i nearby z_i, composed with an arbitrary path from u_i to z_0 on $\mathbb{C} \cup \{+\infty\} \setminus \mathcal{S}$.
- Compute a finite set of generators $\mathcal{X}_i$ for $\mathsf{Expo}_{z_i,\boldsymbol{h}^{z_i}}$ using proposition 6 and add $\Delta_{\gamma_i} \mathrm{X} \Delta_{\gamma_i}^{-1}$ to $\mathcal{M}$ for all $\mathrm{X} \in \mathcal{X}_i$.
- Add $\Delta_{\gamma_i} \Delta_{\circlearrowleft z_i} \Delta_{\gamma_i}^{-1}$ to $\mathcal{M}$.
- For each $\boldsymbol{\alpha} \in \boldsymbol{\mathcal{D}}^{\mathrm{gen}}$ with $\boldsymbol{\mathcal{D}}^{\mathrm{gen}}$ as in remark 13, add $\Delta_{\gamma_i} \Delta_{(z_i,\boldsymbol{k},\boldsymbol{\alpha}^< \to \boldsymbol{\alpha}^>)} \Delta_{\gamma_i}^{-1}$ to $\mathcal{M}$.

return $\mathcal{M}$

THEOREM 14. *With $\mathcal{M}$ constructed as above, the differential Galois group $\mathcal{G}_{L,\boldsymbol{h}^{z_0}}$ is generated by $\mathcal{M}$ as a closed algebraic subgroup of $\mathsf{Mat}_n(\mathbb{C})$.*

PROOF. Assume that $f \in \mathcal{F}\langle h_1^{z_0}, \ldots, h_n^{z_0}\rangle$ is fixed by each element of $\mathcal{M}$. We have to prove that $f \in \mathcal{F}$. Given a singularity z_i, let $\tilde{g}$ be the "continuation" of f along γ_i^{-1} (which involves analytic continuation until u_i followed by "decelero-unsummation"). By proposition 7, we have $\tilde{g} \in \mathbb{C}((z))$. From proposition 12 and remark 13, we next deduce that $\tilde{g}$ is convergent. Indeed, since $\tilde{g} \in \mathbb{C}((z))$, its realization $\hat{g}_i$ in the convolution model with critical time $z_i = z^{1/k_i} = z^{q_i/p_i}$ is a function in $\sqrt[q_i]{\zeta_i}$. Consequently, $\Sigma_{\boldsymbol{k},\boldsymbol{\alpha}^<,\boldsymbol{\alpha}^>} = \Sigma_{\boldsymbol{k},\boldsymbol{\beta}^<,\boldsymbol{\beta}^>}$ whenever $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are equivalent. At this point we have shown that f is meromorphic at z_i. But a function which is meromorphic at all points of the Riemann sphere $\mathbb{C} \cup \{+\infty\}$ is actually a rational function. It follows that $f \in \mathcal{F}$. □

REMARK 15. Slightly less effective versions of theorem 14 are due to Ramis [Ram85, MR91]. Both papers crucially use previous work by Écalle and we followed a similar approach as in the second paper. In the Fuchsian case, i.e. in absence of divergence, the result is due to Schlesinger [Sch95, Sch97].

REMARK 16. We have tried to keep our exposition as short as possible by considering only "directional Stokes-morphisms". In fact, Écalle's theory of resurgent functions gives a more fine-grained control over what happens in the convolution model by considering the *pointed alien derivatives* $\dot{\Delta}_\omega$ for $\omega \in \dot{\mathbb{C}}$. Modulo the identification of functions in the formal model, the convolution models and the geometric model via accelero-summation, the pointed alien derivatives commute with the usual derivation ∂. Consequently, if f is a solution to $L f = 0$, then we also have $L \dot{\Delta}_\omega f = 0$. In particular, given the canonical basis of solutions $\boldsymbol{h}^0$ to $L f = 0$, there exists a unique matrix B_ω with

$$\dot{\Delta}_\omega \boldsymbol{h}^0 = B_\omega \boldsymbol{h}^0.$$

This equation is called the *bridge equation*. Since $\hat{f}_i$ admits only a finite number of singularities and the alien derivations "translate singularities", we have $\dot{\Delta}_\omega^l \boldsymbol{h}^0 = 0$ for some l, so the matrices B_ω are nilpotent. More generally, if $\omega_1, \ldots, \omega_r \in \mathbb{C}^{\neq}$ are $\mathbb{N}$-linearly independent, then all elements in the algebra generated by $B_{\omega_1}, \ldots, B_{\omega_r}$ are nilpotent.

It is easily shown that the Stokes morphisms correspond to the exponentials $\mathrm{e}^{\dot{\Delta}_\theta}$ of directional Alien derivations $\dot{\Delta}_\theta = \sum_{\omega \in \mathrm{e}^{\theta \mathrm{i}} \mathbb{R}^>} \dot{\Delta}_\omega$. This yields a way to reinterpret the Stokes matrices in terms of the B_ω with $\omega \in \mathrm{e}^{\theta \mathrm{i}} \mathbb{R}^>$. In particular, the preceding discussion implies that the Stokes matrices are unipotent. The extra flexibility provided by pointwise over directional alien derivatives admits many applications, such as the preservation of realness [Men96]. For further details, see [É85, É87, É92, É93].

2.6. Effective complex numbers

A complex number z is said to be *effective* if there exists an *approximation algorithm* for z which takes $\varepsilon \in \mathbb{N}^> 2^{\mathbb{Z}}$ on input and which returns an *ε-approximation* $\tilde{z} \in (\mathbb{Z} + \mathrm{i}\,\mathbb{Z})\, 2^{\mathbb{Z}}$ of z for which $|\tilde{z} - z| < \varepsilon$. The *time complexity* of this approximation algorithm is the time $T(d)$ it takes to compute a 2^{-d}-approximation for z. It is not hard to show that the set $\mathbb{C}^{\mathrm{eff}}$ of effective complex numbers forms a field. However, given $z \in \mathbb{C}^{\mathrm{eff}}$ the question whether $z = 0$ is undecidable. The following theorems were proved in [CC90, vdH99, vdH01b].

THEOREM 17. *Let $L \in \mathbb{Q}^{\mathrm{alg}}(z)[\partial]$, $z_0 \in \mathbb{Q}^{\mathrm{alg}} \setminus \mathcal{S}$, $\boldsymbol{v} \in (\mathbb{Q}^{\mathrm{alg}})^n$ and $f = \boldsymbol{v} \cdot \boldsymbol{h}^{z_0}$. Given a broken line path $\gamma = z_0 \to \cdots \to z_k$ on $\mathbb{C}^{\mathrm{eff}} \setminus \mathcal{S}$, we have*

a) *The value $f(\gamma)$ of the analytic continuation of f at the end-point of γ is effective.*
b) *There exists an approximation algorithm of time complexity $O(d \log^3 d \log^2 \log d)$ for $f(\gamma)$, when not counting the approximation time of the input data L, γ and $\boldsymbol{v}$.*
c) *There exists an algorithm which computes an approximation algorithm for $f(\gamma)$ as in (b) as a function of L, γ and $\boldsymbol{v}$.*

THEOREM 18. *Let* $L \in \mathbb{Q}^{\mathrm{alg}}(z)[\partial]$ *be regular singular in* 0. *Let* $z_0 \in \mathbb{Q}^{\mathrm{alg}} \setminus \mathcal{S}$, $\boldsymbol{v} \in (\mathbb{Q}^{\mathrm{alg}})^n$ *and* $f = \boldsymbol{v} \cdot \boldsymbol{h}^{z_0}$. *Then* $f(z)$ *is well-defined for all sufficiently small* γ *on the effective Riemann surface* $\dot{\mathbb{C}}^{\mathrm{eff}}$ *of* log *above* $\mathbb{C}^{\mathrm{eff}}$, *and*

a) $f(\gamma)$ *is effective.*

b) *There exists an approximation algorithm of time complexity* $O(d \log^3 d \log^2 \log d)$ *for* $f(\gamma)$, *when not counting the approximation time of the input data* L, γ *and* $\boldsymbol{v}$.

c) *There exists an algorithm which computes an approximation algorithm for* $f(\gamma)$ *as in (*b*) as a function of* L, γ *and* $\boldsymbol{v}$.

In general, the approximation of $f(\gamma)$ involves the existence of certain bounds. In each of the above theorems, the assertion (c) essentially states that there exists an algorithm for computing these bounds as a function of the input data. This property does not merely follow from (a) and (b) alone.

The following theorem has been proved in [vdH05b].

THEOREM 19. *Let* $L \in \mathbb{Q}^{\mathrm{alg}}(z)[\partial]$ *be singular in* 0. *Let* $\boldsymbol{k}$ *be as in theorem 10 and* $\Gamma = \{(\boldsymbol{k}, \boldsymbol{\alpha}, z) \in \mathrm{dom}_{\mathrm{as}}\, \boldsymbol{h}^0 : \alpha_1, \ldots, \alpha_p, z \in \mathbb{C}^{\mathrm{eff}}\}$. *Given* $f = \boldsymbol{v} \cdot \boldsymbol{h}^0$ *with* $\boldsymbol{v} \in (\mathbb{C}^{\mathrm{eff}})^n$ *and* $\gamma \in \Gamma$, *we have*

a) $f(\gamma)$ *is effective.*

b) *There exists an approximation algorithm of time complexity* $O(d \log^4 d \log \log d)$ *for* $f(\gamma)$, *when not counting the approximation time of the input data* L, γ *and* $\boldsymbol{v}$.

c) *There exists an algorithm which computes an approximation algorithm for* $f(\gamma)$ *as in (*b*) as a function of* L, γ *and* $\boldsymbol{v}$.

If we replace $\mathbb{Q}^{\mathrm{alg}}$ by an arbitrary effective algebraically closed subfield $\mathbb{K}$ of $\mathbb{C}^{\mathrm{eff}}$, then the assertions (a) and (c) in the three above theorems remain valid (see [vdH05a, vdH03] in the cases of theorems 17 and 18), but the complexity in (b) usually drops back to $O(d^2 \log^{O(1)} d)$. Notice also that we may replace $f(\gamma)$ by the transition matrix along γ in each of the theorems. The following theorem summarizes the results from section 2.

THEOREM 20. *Let* $\mathbb{K}$ *be an effective algebraically closed constant field of* $\mathbb{C}^{\mathrm{eff}}$. *Then there exists an algorithm which takes* $L \in \mathbb{K}(z)[\partial]$ *and* $z_0 \in \mathbb{K} \cup \{\infty\}$ *on input, and which computes a finite set* $\mathcal{M} \subseteq \mathsf{Mat}(\mathbb{C}^{\mathrm{eff}})$, *such that*

a) *The group* $\mathcal{G}_{L, \boldsymbol{h}^{z_0}}$ *is generated by* $\mathcal{M}$ *as a closed algebraic subgroup of* $\mathsf{Mat}_n(\mathbb{C})$.

b) *If* $\mathbb{K} = \mathbb{Q}^{\mathrm{alg}}$, *then the entries of the matrices in* $\mathcal{M}$ *have time complexity* $O(d \log^4 d \log \log d)$.

PROOF. It is classical that the set $\mathbb{C}^{\mathrm{eff}}$ of effective complex numbers forms a field. Similarly, the set $\mathbb{C}^{\mathrm{fast}}$ of effective complex numbers with an approximation algorithm of time complexity $O(d \log^4 d \log \log d)$ forms a field, since the operations $+$, $-$, $\times$ and $/$ can all be performed in time $O(d \log d \log \log d)$. In particular, the classes of matrices with entries in $\mathbb{C}^{\mathrm{eff}}$ resp. $\mathbb{C}^{\mathrm{fast}}$ are stable under the same operations. Now in the algorithm `Compute_generators`, we may take broken-line paths with vertices above $\mathbb{K}$ for the γ_i. Hence (a) and (b) follow from theorem 19(a) resp. (b) and the above observations. □

Given $\varepsilon \in \mathbb{N}^{>} 2^{\mathbb{Z}}$, we may endow $\mathbb{C}^{\mathrm{eff}}$ with an approximate zero-test for which $z = 0$ if and only if $|z| < \varepsilon$. We will denote this field by $\mathbb{C}^{\approx \varepsilon}$. Clearly, this zero-test is not compatible with the field structure of $\mathbb{C}^{\mathrm{eff}}$. Nevertheless, any finite computation, which can be carried out in $\mathbb{C}^{\mathrm{eff}}$ with an oracle for zero-testing, can be carried out in exactly the same way in $\mathbb{C}^{\approx \varepsilon}$ for a sufficiently small ε. Given $z \in \mathbb{C}^{\mathrm{eff}}$, we will denote by $z^{\approx \varepsilon} \in \mathbb{C}^{\approx \varepsilon}$ the "cast" of z to $\mathbb{C}^{\approx \varepsilon}$ and similarly for matrices with coefficients in $\mathbb{C}^{\mathrm{eff}}$.

REMARK 21. In practice [vdH04], effective complex numbers z usually come with a natural bound $M \in \mathbb{N}^> 2^{\mathbb{Z}}$ for $|z|$. Then, given $\varepsilon \in \mathbb{N}^> 2^{\mathbb{Z}}$ with $\varepsilon < 1$, it is even better to use the approximate zero-test $z = 0$ if and only if $|z| < \varepsilon\, M$. Notice that the bound M usually depends on the internal representation of z and not merely on z as a number in $\mathbb{C}^{\text{eff}}$.

3. FACTORING LINEAR DIFFERENTIAL OPERATORS

Let $\mathbb{K}$ be an effective algebraically closed subfield of $\mathbb{C}^{\text{eff}}$. Consider a monic linear differential operator $L = \partial^n + L_{n-1}\,\partial + \cdots + L_0 \in \mathcal{F}[\partial]$, where $\mathcal{F} = \mathbb{K}(z)$. In this section, we present an algorithm for finding a non-trivial factorization $L = K_1\,K_2$ with $K_1, K_2 \in \mathcal{F}[\partial]$ whenever such a factorization exists.

3.1. Factoring L and invariant subspaces under $\mathcal{G}_{L,\boldsymbol{h}}$

Let $\boldsymbol{h} = (h_1, \ldots, h_n) \in \mathcal{K}^n$ be a basis of solutions for the equation $L\,f = 0$, where $\mathcal{K} \supseteq \mathcal{F}$ is an abstract differential field. We denote the Wronskian of $\boldsymbol{h}$ by

$$W_{\boldsymbol{h}} = W_{h_1,\ldots,h_n} = \begin{vmatrix} h_1 & \cdots & h_n \\ \vdots & & \vdots \\ h_1^{(n-1)} & \cdots & h_n^{(n-1)} \end{vmatrix}.$$

It is classical (and easy to check) that

$$L\,f = \frac{W_{f,h_1,\ldots,h_n}}{W_{h_1,\ldots,h_n}}. \tag{7}$$

When expanding the determinant $W_{f,h_1,\ldots,h_n}$ in terms of the matrices

$$W_i = W_{\boldsymbol{h},i} = \begin{vmatrix} h_1 & \cdots & h_n \\ \vdots & & \vdots \\ h_1^{(n-i-1)} & \cdots & h_n^{(n-i-1)} \\ h_1^{(n-i+1)} & \cdots & h_n^{(n-i+1)} \\ \vdots & & \vdots \\ h_1^{(n-1)} & \cdots & h_n^{(n-1)} \end{vmatrix},$$

it follows that

$$L = \partial^n - \frac{W_1}{W_0}\,\partial^{n-1} + \cdots + (-1)^n\,\frac{W_n}{W_0}.$$

Denoting by $\varphi^\dagger$ the logarithmic derivative of φ, it can also be checked by induction that

$$\tilde{L} = \left(\partial - \left(\frac{W_{h_1,\ldots,h_n}}{W_{h_1,\ldots,h_{n-1}}}\right)^\dagger\right) \cdots \left(\partial - \left(\frac{W_{h_1,h_2}}{W_{h_1}}\right)^\dagger\right)(\partial - W_{h_1}^\dagger)$$

admits $h_1, \ldots, h_n$ as solutions, whence $\tilde{L} = L$, using Euclidean division in the skew polynomial ring $\mathcal{F}[\partial]$.

PROPOSITION 22.

a) *If L admits a factorization $L = K_1\,K_2$, then $\mathcal{G}_{L,\boldsymbol{h}}$ leaves $\ker K_2$ invariant.*

b) *If V is an invariant subvector space of $\mathcal{G}_{L,\boldsymbol{h}}$, then L admits a factorization $L = K_1\,K_2$ with $V = \ker K_2$.*

PROOF. Assume that L admits a factorization $L = K_1 K_2$. Then, given $f \in \ker K_2$ and $M \in \mathcal{G}_{L,\boldsymbol{h}}$, we have $K_2 f = 0 = M K_2 f = K_2 M f$, whence $M f \in \ker K_2$. Conversely, assume that V is an invariant subvector space of $\mathcal{G}_{L,\boldsymbol{h}}$ and let $\boldsymbol{g}$ be a basis of V. Then we observe that $M W_{\boldsymbol{g},i} = (\det M) W_{\boldsymbol{g},i}$ for all i. Consequently,

$$M \frac{W_{\boldsymbol{g},i}}{W_{\boldsymbol{g},0}} = \frac{W_{\boldsymbol{g},i}}{W_{\boldsymbol{g},0}}$$

for all i, so that $W_{\boldsymbol{g},i}/W_{\boldsymbol{g},0} \in \mathcal{F}$, by corollary 3. Hence

$$K_2 = \partial^r - \frac{W_{\boldsymbol{g},1}}{W_{\boldsymbol{g},0}} \partial^{r-1} + \cdots + (-1)^r \frac{W_{\boldsymbol{g},r}}{W_{\boldsymbol{g},0}} \tag{8}$$

is a differential operator with coefficients in $\mathcal{F}$ which vanishes on V. But this is only possible if K_2 divides L. □

3.2. A lemma from linear algebra

LEMMA 23. *Let $\mathcal{A}$ be a non-unitary algebra of nilpotent matrices in $\mathsf{Mat}_n(\mathbb{K})$. Then there exists a basis of $\mathbb{K}^n$ in which M is lower triangular for all $M \in \mathcal{A}$.*

PROOF. Let $M \in \mathcal{A}$ be a matrix such that $V = \operatorname{im} M$ is a non-zero vector space of minimal dimension. Given $\boldsymbol{v} \in V$ and $N \in \mathcal{A}$, we claim that $N\boldsymbol{v} \in \ker M$. Assume the contrary, so that $0 \neq MN\boldsymbol{v} \in \operatorname{im} MN \subseteq V$. By the minimality hypothesis, we must have $\operatorname{im} MN = V$. In particular, $\boldsymbol{v} \in \operatorname{im} MN$ and $0 \neq MN\boldsymbol{v} \in \operatorname{im} MNMN$. Again by the minimality hypothesis, it follows that $\operatorname{im} MNMN = V$. In other words, the restriction of MN to V is an isomorphism on V. Hence MN admits a non-zero eigenvector in V, which contradicts the fact that MN is nilpotent.

Let us now prove the lemma by induction over n. If $n \leqslant 1$ or $\mathcal{A} = 0$, then we have nothing to do, so assume that $n > 1$ and $\mathcal{A} \neq 0$. We claim that $\mathbb{K}^n$ admits a non-trivial invariant subvector space W. Indeed, we may take $W = V$ if $\mathcal{A} V = 0$ and $W = \mathcal{A} V$ if $\mathcal{A} V \neq 0$. Now consider a basis $(\boldsymbol{b}_{m+1}, \ldots, \boldsymbol{b}_n)$ of W and complete it to a basis $(\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n)$ of $\mathbb{K}^n$. Then each matrix in $\mathcal{A}$ is lower triangular with respect to this basis. Let $\mathcal{A}_1$ and $\mathcal{A}_2$ be the algebras of lower dimensional matrices which occur as upper left resp. lower right blocks of matrices in $\mathcal{A}$. We conclude by applying the induction hypothesis on $\mathcal{A}_1$ and $\mathcal{A}_2$. □

Let $\mathcal{M}$ be a finite set of non-zero nilpotent matrices. If all matrices in the $\mathbb{K}$-algebra $\mathcal{A}$ generated by $\mathcal{M}$ are nilpotent, then it is easy to compute a basis for which all matrices in $\mathcal{M}$ are lower triangular. Indeed, setting $K_i = \bigcap_{M \in \mathcal{M}^i} \ker M$ for all i, we first compute a basis of K_1. We successively complete this basis into a basis of K_2, K_3 and so on until $K_p = \mathbb{K}^n$.

If not all matrices in $\mathcal{A}$ are nilpotent, then the proof of lemma 23 indicates a method for the computation of a matrix in $\mathcal{A}$ which is not nilpotent. Indeed, we start by picking an $M \in \mathcal{M}$ and set $V := \operatorname{im} M^{p-1}$, where p is smallest with $M^p = 0$. We next set $\mathcal{N} := \mathcal{M} \setminus \{M\}$ and iterate the following loop. Take a matrix $N \in \mathcal{N}$ and distinguish the following three cases:

$\boldsymbol{MNV = 0}$**.** Set $\mathcal{N} := \mathcal{N} \setminus \{N\}$ and continue.

$\boldsymbol{0 \subsetneq MNV \subsetneq V}$**.** Set $M := MNM$, $V := \operatorname{im} M$ and continue.

$\boldsymbol{MNV = V}$**.** Return the non-nilpotent matrix MN.

At the end of our loop, we either found a non-nilpotent matrix, or we have $NV \subseteq \ker M$ for all $N \in \mathcal{M}$. In the second case, we obtain a non-trivial invariant subspace of $\mathbb{K}^n$ as in the proof of lemma 23 and we recursively apply the algorithm on this subspace and a complement. In fact, the returned matrix is not even *monopotent* (i.e. not of the form $\lambda + N$, where N is a nilpotent matrix), since it both admits zero and a non-zero number as eigenvalues.

3.3. Computation of non-trivial invariant subspaces

Proposition 22 in combination with theorem 20 implies that the factorization of linear differential operators in $\mathcal{F}[\partial]$ reduces to the computation of non-trivial invariant subvector spaces under the action of $\mathcal{M}_{L,\boldsymbol{h}}$ whenever they exist.

In this section, we will first solve a slightly simpler problem: assuming that $\mathbb{K}$ is an effective algebraically closed field and given a finite set of matrices $\mathcal{M} \subseteq \mathsf{Mat}_n(\mathbb{K})$, we will show how to compute a non-trivial invariant subspace V of $\mathbb{K}^n$ under the action of $\mathcal{M}$, whenever such a V exists.

Good candidate vectors $\boldsymbol{v}$. Given a vector $\boldsymbol{v} \in \mathbb{K}^n$ it is easy to compute the smallest subspace $\mathsf{Inv}_{\mathcal{M}}(\boldsymbol{v})$ of $\mathbb{K}^n$ which is invariant under the action of $\mathcal{M}$ and which contains $\boldsymbol{v}$. Indeed, starting with a basis $\mathcal{B} = \{\boldsymbol{v}\}$, we keep enlarging $\mathcal{B}$ with elements in $\mathcal{M}\ \mathcal{B} \setminus \mathsf{Vect}(\mathcal{B})$ until saturation. Since $\mathcal{B}$ will never contain more than n elements, this algorithm terminates. A *candidate vector* $\boldsymbol{v} \in \mathbb{K}^n$ for generating a non-trivial invariant subspace of $\mathbb{K}^n$ is said to be *good* if $0 < \dim \mathsf{Inv}_{\mathcal{M}}(\boldsymbol{v}) < n$.

The $\mathbb{K}$-algebra generated by $\mathcal{M}$. We notice that $V \subseteq \mathbb{K}^n$ is an invariant subspace for $\mathcal{M}$, if and only if V is an invariant subspace for the $\mathbb{K}$-algebra $\mathsf{Alg}(\mathcal{M})$ generated by $\mathcal{M}$. Again it is easy to compute a basis for $\mathsf{Alg}(\mathcal{M})$. We start with a basis $\mathcal{B}$ of $\mathsf{Vect}(\mathcal{M})$ and keep adjoining elements in $\mathcal{B}^2 \setminus \mathsf{Vect}(\mathcal{B})$ to $\mathcal{B}$ until saturation. We will avoid the explicit basis of $\mathsf{Alg}(\mathcal{M})$, which may contain as much as n^2 elements, and rather focus on the efficient computation of good candidate vectors.

$\mathcal{M}$-splittings. A decomposition $\mathbb{K}^n = E_1 \oplus \cdots \oplus E_k$, where $E_1, \ldots, E_k$ are non-empty vector spaces, will be called an *$\mathcal{M}$-splitting* of $\mathbb{K}^n$, if the projections $P_i = P_{E_i}$ of $\mathbb{K}^n$ on E_i are all in $\mathsf{Alg}(\mathcal{M})$. Then, given $M \in \mathsf{Mat}_n(\mathbb{K})$, we have $M \in \mathsf{Alg}(\mathcal{M})$ if and only if $P_i M P_j \in \mathsf{Alg}(\mathcal{M})$ for all i, j. If we choose a basis for $\mathbb{K}^n$ which is a union of bases for the E_i, we notice that the $P_i M P_j$ are $\dim E_i \times \dim E_j$ block matrices. In the above algorithm for computing the $\mathbb{K}$-algebra generated by $\mathcal{M}$ it now suffices to compute with block matrices of this form. In particular, the computed basis of $\mathsf{Alg}(\mathcal{M})$ will consist of such matrices. The trivial decomposition $\mathbb{K}^n = \mathbb{K}^n$ is clearly an $\mathcal{M}$-splitting. Given $N \in \mathsf{Alg}(\mathcal{M})$, we notice that any $\{N\}$-splitting is also an $\mathcal{M}$-splitting.

Refining $\mathcal{M}$-splittings. An $\mathcal{M}$-splitting $\mathbb{K}^n = F_1 \oplus \cdots \oplus F_l$ is said to be *finer* than the $\mathcal{M}$-splitting $\mathbb{K}^n = E_1 \oplus \cdots \oplus E_k$ if E_i is a direct sum of a subset of the F_j for each i. Given an $\mathcal{M}$-splitting $\mathbb{K}^n = F_1 \oplus \cdots \oplus F_l$ and an arbitrary element $M \in \mathsf{Alg}(\mathcal{M})$, we may obtain a finer *$\mathcal{M}$-splitting w.r.t M* as follows. Let $i \in \{1, \ldots, k\}$ and consider $M_i = P_i M P_i$. If $\lambda_1, \ldots, \lambda_p$ are the eigenvalues of M_i, then $E_i = \ker (M_i - \lambda_1)^{n_i} \oplus \cdots \oplus \ker (M_i - \lambda_k)^{n_i}$ is an $(P_i \mathcal{M} P_i)$-splitting of E_i, where $n_i = \dim E_i$. Collecting these $(P_i \mathcal{M} P_i)$-splittings, we obtain a finer $\mathcal{M}$-splitting $F_1 \oplus \cdots \oplus F_l$ of $\mathbb{K}^n$. This $\mathcal{M}$-splitting, which is said to be *refined w.r.t. M*, has the property that $P_{F_i} M P_{F_i}$ is monopotent on F_i for each i, with unique eigenvalue λ_{M,F_i}.

We now have the following algorithm for computing non-trivial $\mathcal{M}$-invariant subspaces of $\mathbb{K}^n$ when they exist.

Algorithm `Invariant_subspace`$(\mathcal{M})$
Input: a set of non-zero matrices in $\mathsf{Mat}_n(\mathbb{K})$
Output: an $\mathcal{M}$-invariant subspace of $\mathbb{K}^n$ or **fail**

Step 1. [Initial $\mathcal{M}$-splitting]
Compute a "random non-zero element" N of $\mathsf{Alg}(\mathcal{M})$
Compute an $\mathcal{M}$-splitting $\mathbb{K}^n = E_1 \oplus \cdots \oplus E_k$ w.r.t. N and each $M \in \mathcal{M}$
$\mathcal{D} := \varnothing$

Step 2. [One dimensional components]
For every E_i with $\dim E_i = 1$ and $E_i \notin \mathcal{D}$, do the following:
Pick a $\boldsymbol{v} \in E_i^{\neq}$ and compute $\mathsf{Inv}_{\mathcal{M}}(\boldsymbol{v})$
If $\mathsf{Inv}_{\mathcal{M}}(\boldsymbol{v}) \subsetneq \mathbb{K}^n$ then return $\mathsf{Inv}_{\mathcal{M}}(\boldsymbol{v})$
Otherwise, set $\mathcal{D} := \mathcal{D} \cup E_i$
If $\dim E_i = 1$ for all i then return **fail**

Step 3. [Higher dimensional components]
Let i be such that $\dim E_i > 1$
Let $\mathcal{M}_i := \{P_i (M - \lambda_{M,E_i}) P_i : M \in \mathcal{M}\}$
Let $K_i := E_i \cap \bigcap_{M \in \mathcal{M}_i} \ker M$
If $K_i = 0$ then go to step 4 and otherwise to step 5

Step 4. [Non-triangular case]
Let $N \in \mathsf{Alg}(\mathcal{M}_i)$ be non-monopotent on E_i (cf. previous section)
Refine the $\mathcal{M}$-splitting w.r.t. N and return to step 2

Step 5. [Potentially triangular case]
Choose $\boldsymbol{v} \in K_i$ and compute $\mathsf{Inv}_{\mathcal{M}}(\boldsymbol{v})$
If $\mathsf{Inv}_{\mathcal{M}}(\boldsymbol{v}) \subsetneq \mathbb{K}^n$ then return $\mathsf{Inv}_{\mathcal{M}}(\boldsymbol{v})$
Otherwise, let N be in $\mathsf{Alg}(\mathcal{M})$ with $P_i N \boldsymbol{v} \notin K_i$
Refine the $\mathcal{M}$-splitting w.r.t. N
If this yields a finer $\mathcal{M}$-splitting then return to step 2
Otherwise, set $\mathcal{M} := \mathcal{M} \cup \{N\}$ and repeat step 5

The algorithm needs a few additional explanations. In step 1, we may take N to be an arbitrary element in $\mathcal{M}$. However, it is better to take a "small random expression in the elements of $\mathcal{M}$" for N. With high probability, this yields an $\mathcal{M}$-splitting which will not need to be refined in the sequel. Indeed, the subset of matrices in $\mathsf{Alg}(\mathcal{M})$ which yield non-maximal $\mathcal{M}$-splittings is a closed algebraic subset of measure zero, since it is determined by coinciding eigenvalues. In particular, given an $\mathcal{M}$-splitting $\mathbb{K}^n = E_1 \oplus \cdots \oplus E_k$ w.r.t. N, it will usually suffice to check that each $M \in \mathcal{M}$ is monopotent on each E_i, in order to obtain an $\mathcal{M}$-splitting w.r.t. the other elements in $\mathcal{M}$.

Throughout the algorithm, the $\mathcal{M}$-splitting gets finer and finer, so the $\mathcal{M}$-splitting ultimately remains constant. From this point on, the space K_i can only strictly decrease in step 5, so K_i also remains constant, ultimately. But then we either find a non-trivial invariant subspace in step 5, or all components of the $\mathcal{M}$-splitting become one-dimensional. In the latter case, we either obtain a non-trivial invariant subspace in step 1, or a proof that $\mathsf{Inv}_{\mathcal{M}}(\boldsymbol{v}) = \mathbb{K}^n$ for every $\boldsymbol{v} \in E_1^{\neq} \cup \cdots \cup E_n^{\neq}$ (and thus for every $\boldsymbol{v} \in \mathbb{K}^n \setminus 0$).

REMARK 24. Assume that $\mathbb{K}$ is no longer an effective algebraically closed field, but rather a field $\mathbb{C}^{\approx\varepsilon}$ with an approximate zero-test. In that case, we recall that a number which is approximately zero is not necessarily zero. On the other hand, a number which is not approximately zero is *surely* non-zero. Consequently, in our algorithm for the computation of $\mathsf{Inv}(\boldsymbol{v})$, the dimension of $\mathsf{Inv}(\boldsymbol{v})$ can be too small, but it is never too large. In particular, if the algorithm `Invariant_subspace` fails, then the approximate proof that $\mathsf{Inv}_{\mathcal{M}}(\boldsymbol{v}) = \mathbb{K}^n$ for every $\boldsymbol{v} \in E_1^{\neq} \cup \cdots \cup E_n^{\neq}$ yields a genuine proof that there are no non-trivial invariant subspaces.

3.4. Factoring linear differential operators

Putting together the results from the previous sections, we now have the following algorithm for finding a right factor of L.

Algorithm `Right_factor`(L)
Input: $L = \partial^n + L_{n-1}\,\partial^{n-1} + \cdots + L_0 \in \mathbb{K}(z)[\partial]$
Output: a non-trivial right-factor of L or **fail**

Step 1. [Compute generators]
Choose $z_0 \in \mathbb{K} \setminus \mathcal{S}$ and let $\boldsymbol{h} = \boldsymbol{h}^{z_0}$
Compute a finite set $\mathcal{M} \subseteq \mathsf{GL}_n(\mathbb{C}^{\text{eff}})$ of generators for $\mathcal{G}_{L,\boldsymbol{h}}$ (cf. theorem 20)

Step 2. [Initial precision]
$T := \max(\deg L_0, \ldots, \deg L_{n-1}) + 1$
$\delta := 2^{-32}$
while $\delta' := \min\{M_{i,j}/M_{i',j'} \colon M \in \mathcal{M}, M_{i',j'}^{\approx \delta/2^T} \neq 0\} < \delta$ **do** $\delta := \delta'$
$\varepsilon := \delta/2^T$

Step 3. [Produce invariant subspace]
Let $V :=$ `Invariant_subspace`$(\mathcal{M}^{\approx\varepsilon})$
If $V =$ **fail** then return **fail**
Let $B \in \mathsf{Mat}_{n,r}(\mathbb{C}^{\approx\varepsilon})$ be a column basis of V
Let $\boldsymbol{g} := {}^t B\,\boldsymbol{h} \in ([\mathbb{C}^{\text{eff}}]^{\approx\varepsilon}[[z]]^{\text{eff}})^r$

Step 4. [Produce and check guess]
Let $K := \partial^r - \frac{W_{\boldsymbol{g},1}}{W_{\boldsymbol{g},0}}\,\partial^{r-1} + \cdots + (-1)^r\,\frac{W_{\boldsymbol{g},r}}{W_{\boldsymbol{g},0}}$
Divide L by K, producing $Q, R \in [\mathbb{C}^{\approx\varepsilon}((z))]^{\text{eff}}[\partial]$ with $L = Q\,K + R$
If $R \neq 0 \bmod z^T$ then go to step 5
Reconstruct $\tilde{Q}, \tilde{K} \in \mathbb{K}(z)[\partial]$ from Q and K with precision (ε, T)
If we obtain no good approximations or $L \neq \tilde{Q}\,\tilde{K}$ then go to step 5
Return $\tilde{K}$

Step 5. [Increase precision]
$T := 2\,T$
$\varepsilon := \delta/2^T$
Go to step 3

The main idea behind the algorithm is to use proposition 22 in combination with `Invariant_subspace` so as to provide good candidate right factors of L in $\mathbb{C}^{\text{eff}}((z))[\partial]$. Using reconstruction of coefficients in $\mathbb{K}(z)$ from Laurent series in $\mathbb{C}^{\text{eff}}((z))$ with increasing precisions, we next produce good candidate right factors in $\mathbb{K}(z)$. We keep increasing the precision until we find a right factor or a proof that L is irreducible. Let us detail the different steps a bit more:

Step 2. We will work with power series approximations of T terms and approximate zero-tests in $\mathbb{C}^{\approx\varepsilon}$. The degree of a rational function P/Q is defined by $\deg P/Q = \max(\deg P, \deg Q)$. The initial precisions T and $-\log \varepsilon$ have been chosen as small as possible. Indeed, we want to take advantage of a possible quick answer when computing with a small precision (see also the explanations below of step 5).

Step 3. If `Invariant_subspace` fails, then there exists no factorization of L, by remark 24. Effective power series and Laurent series are defined in a similar way as effective real numbers (in particular, we don't assume the existence of an effective zero-test). Efficient algorithm for such computations are described in [vdH02].

Step 4. The reconstruction of $\tilde{Q}$ and $\tilde{K}$ from Q and K contains two ingredients: we use Padé approximation to find rational function approximations of degree $\leqslant T$ and the LLL-algorithm to approximate numbers $\mathbb{C}^{\approx\varepsilon}$ by numbers in $\mathbb{K}$.

Step 5. Doubling the precision at successive steps heuristically causes the computation time to increase geometrically at each step. In particular, unsuccessful computations at lower precisions don't take much time with respect to the last successful computation with respect to the required precision. Instead of multiplying the precisions by two, we also notice that it would be even better to increase by a factor which doubles the estimated computation time at each step. Of course, this would require a more precise complexity analysis of the algorithm.

The problem of reconstructing elements in $\mathbb{K}$ from elements in $\mathbb{C}^{\approx\varepsilon}$ is an interesting topic on its own. In theory, one may consider the polynomial algebra over $\mathbb{Z}$ generated by all coefficients occurring in L and the number z we wish to reconstruct. We may then apply the LLL-algorithm [LLL82] on the lattice spanned by $\sqrt{T}$ monomials of smallest total degree (for instance) and search for minimal $\leqslant \sqrt{T}$-digit relations. If $\mathbb{K} = \mathbb{Q}^{\mathrm{alg}}$ is the algebraic closure of $\mathbb{Q}$, then we may simply use the lattice spanned by the first n powers of z.

At a sufficiently large precision T, the LLL-algorithm will ultimately succeed for all coefficients of a candidate factorization which need to be reconstructed. If there are no factorizations, then the algorithm will ultimately fail at step 3. This proofs the termination of `Right_factor`.

REMARK 25. In practice, and especially if $\mathbb{K} \neq \mathbb{Q}^{\mathrm{alg}}$, it would be nice to use more of the structure of the original problem. For instance, a factorization of L actually yields relations on the coefficients which we may try to use. For high precision computations, it is also recommended to speed the LLL-algorithm up using a similar dichotomic algorithm as for fast g.c.d. computations [Moe73, PW02].

REMARK 26. Notice that we did not use bounds for the degrees of coefficients of possible factors in our algorithm. If a bound T^B is available, using techniques from [BB85, vH97, vdPS03], then one may take $T := \min(2T, T^B)$ instead of $T := 2T$ in step 5. Of course, bounds for the required precision ε are even harder to obtain. See [BB85] for some results in that direction.

4. COMPUTING DIFFERENTIAL GALOIS GROUPS

4.1. Introduction

Throughout this section, $\mathbb{F}$ will stand for the field $\mathbb{C}^{\approx\varepsilon}$ of effective complex number with the approximate zero-test at precision $\varepsilon > 0$. This field has the following properties:

EH1. We have an effective zero-test in $\mathbb{F}$.

EH2. There exists an algorithm which takes on input $\boldsymbol{c} \in (\mathbb{F}^{\neq})^n$ and which computes a finite set of generators for the $\mathbb{Z}$-vector space of integers $\boldsymbol{k} \in \mathbb{Z}^n$ with $\boldsymbol{c}^{\boldsymbol{k}} = 1$.

EH3. There exists an algorithm which takes on input $\boldsymbol{c} \in \mathbb{F}^n$ and which computes a finite set of generators for the $\mathbb{Z}$-vector space of integers $\boldsymbol{k} \in \mathbb{Z}^n$ with $\boldsymbol{c} \cdot \boldsymbol{k} = 0$.

EH4. $\mathbb{F}$ is closed under exponentiation and logarithm.

Indeed, we obtain **EH2** and **EH3** using the LLL-algorithm. Some of the results in this section go through when only a subset of the conditions are satisfied. In that case, we notice that **EH2** $\Rightarrow$ **EH1**, **EH3** $\Rightarrow$ **EH1** and **EH4** $\Rightarrow$ (**EH2** $\Leftrightarrow$ **EH3**).

Given a finite set of matrices $\mathcal{M} \subseteq \mathsf{GL}_n(\mathbb{F})$, we give a numerical algorithm for the computation of the smallest closed algebraic subgroup $\mathcal{G} = \langle \mathcal{M} \rangle$ of $\mathsf{GL}_n(\mathbb{F})$ which contains $\mathcal{M}$. We will represent $\mathcal{G}$ by a finite set $\mathcal{F} \subseteq \mathsf{GL}_n(\mathbb{F})$ and the finite basis $\mathcal{B} \subseteq \mathsf{GL}(\mathbb{F})$ of a Lie algebra $\mathcal{L}$ over $\mathbb{C}$, such that

$$\mathcal{G} = \mathcal{F}\, \mathrm{e}^{\mathcal{L}},$$

and each $N \in \mathcal{F}$ corresponds to a unique connected component $N\, \mathrm{e}^{\mathcal{L}} = \mathrm{e}^{\mathcal{L}} N$ of $\mathcal{G}$. We will also prove that there exists a precision ε_0 such that the algorithm yields the theoretically correct result for all $\varepsilon < \varepsilon_0$.

4.2. The algebraic group generated by a diagonal matrix

Let $\mathsf{Tor}_n(\mathbb{F})$ be the group of invertible diagonal matrices. Each matrix M has the form $M = \mathrm{Diag}(\boldsymbol{\alpha})$, where $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n)$ is the vector in $(\mathbb{F}^{\neq})^n$ of the elements on the diagonal of M. The coordinate ring $\mathcal{R}$ of $\mathsf{Tor}_n(\mathbb{F})$ is the set $\mathbb{F}[\boldsymbol{\alpha}, \boldsymbol{\alpha}^{-1}]$ of Laurent polynomials in $\boldsymbol{\alpha}$.

Now consider the case when $\mathcal{M}$ consists of a single diagonal matrix $M = \mathrm{Diag}(\boldsymbol{\lambda})$. Let $\mathfrak{i} \subseteq \mathcal{R}$ be the ideal which defines $\langle M \rangle \subseteq \mathsf{Tor}_n(\mathbb{F})$. Given a relation $\boldsymbol{\lambda}^{\boldsymbol{k}} = 1$ $(\boldsymbol{k} \in \mathbb{Z}^n)$ between the λ_i, any power $M^i = \mathrm{Diag}(\boldsymbol{\lambda}^i)$ satisfies the same relation $(\boldsymbol{\lambda}^i)^{\boldsymbol{k}} = 1$, whence $\boldsymbol{\alpha}^{\boldsymbol{k}} - 1 \in \mathfrak{i}$. Let $\mathfrak{j}$ be the ideal generated by all $\boldsymbol{\alpha}^{\boldsymbol{k}} - 1$, such that $\boldsymbol{\lambda}^{\boldsymbol{k}} = 1$.

LEMMA 27. *We have* $\mathfrak{j} = \mathfrak{i}$.

PROOF. We already observed that $\mathfrak{j} \subseteq \mathfrak{i}$. Assuming for contradiction that $\mathfrak{j} \neq \mathfrak{i}$, choose

$$f = \sum_{i=1}^{r} f_i\, \boldsymbol{\alpha}^{\boldsymbol{k}_i} \in \mathfrak{i} \setminus \mathfrak{j}$$

such that r is minimal. If $i \neq j$, then $\boldsymbol{\lambda}^{\boldsymbol{k}_i - \boldsymbol{k}_j} \neq 1$, since otherwise $\boldsymbol{\lambda}^{\boldsymbol{k}_i - \boldsymbol{k}_j} \in \mathfrak{j}$ and $f - f_i\ \boldsymbol{\alpha}^{\boldsymbol{k}_i} + f_i\ \boldsymbol{\alpha}^{\boldsymbol{k}_j} \in \mathfrak{i} \setminus \mathfrak{j}$ has less than r terms. In particular, the vectors $(1, \boldsymbol{\lambda}^{\boldsymbol{k}_i}, \ldots, \boldsymbol{\lambda}^{(r-1)\boldsymbol{k}_i})$ with $i \in \{1, \ldots, r\}$ are linearly independent. But this contradicts the fact that $f(\boldsymbol{\lambda}^j) = \sum_{i=1}^{r} f_i\, \boldsymbol{\lambda}^{j\boldsymbol{k}_i} = 0$ for all $j \in \{0, \ldots, r-1\}$. □

By **EH2**, we may compute a minimal finite set $\boldsymbol{g}_1, \ldots, \boldsymbol{g}_p$ of generators for the $\mathbb{Z}$-vector space of $\boldsymbol{k} \in \mathbb{Z}^n$ with $\boldsymbol{\lambda}^{\boldsymbol{k}} = 1$. We may also compute a basis $\mathcal{B}$ for $\ker \varphi$, where $\varphi \colon \mathbb{Z}^n \to \mathbb{Z}^p$; $\boldsymbol{k} \mapsto (\boldsymbol{k} \cdot \boldsymbol{g}_1, \ldots, \boldsymbol{k} \cdot \boldsymbol{g}_p)$. Then $\mathrm{e}^{\mathcal{L}} = \mathrm{e}^{\mathsf{Vect}(\mathcal{B})}$ is the connected component of $\langle M \rangle$, since $(\mathrm{e}^{\mathcal{L}})^{\boldsymbol{g}_i} = 1$ for all i, and $\mathcal{L}$ cannot be further enlarged while conserving this property.

Let $\mathcal{V} = (\boldsymbol{g}_1\, \mathbb{Q} \oplus \cdots \oplus \boldsymbol{g}_p\, \mathbb{Q}) \cap \mathbb{Z}^n$. We construct a basis $\boldsymbol{h}_1, \ldots, \boldsymbol{h}_n$ of $\mathbb{Z}^n$, by taking $\boldsymbol{h}_i$ to be shortest in $\mathcal{V}$ (if $i \leqslant p$) or $\mathbb{Z}^n$ (if $i > p$), such that $\boldsymbol{h}_i \notin \mathsf{Vect}(\boldsymbol{h}_1, \ldots, \boldsymbol{h}_{i-1})$. This basis determines a toric change of coordinates $\boldsymbol{\alpha} \to \boldsymbol{\alpha}^P$ with $P \in \mathsf{GL}_n(\mathbb{Z})$ such that $\boldsymbol{g}_1, \ldots, \boldsymbol{g}_p \in \mathbb{Z}^p \times 0^{n-p}$ with respect to the new coordinates. Similarly, we may construct a basis $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_p$ of $\mathbb{Z}^p$, by taking each $\boldsymbol{b}_i$ to be shortest in $\mathbb{Z}^p \setminus \mathsf{Vect}(\boldsymbol{b}_1, \ldots, \boldsymbol{b}_{i-1})$ such that $r_i = \min \{r \in \mathbb{N}^{>} \colon r\, \boldsymbol{b}_i \in \mathbb{Z}\, \boldsymbol{g}_1 \oplus \cdots \oplus \mathbb{Z}\, \boldsymbol{g}_p\}$ is maximal. This basis determines a second toric change of coordinates $\boldsymbol{\alpha} \to \boldsymbol{\alpha}^Q$ with $Q \in \mathsf{GL}_n(\mathbb{Z})$ such that $\boldsymbol{g}_i = r_i\, \boldsymbol{e}_i$ $(i = 1, \ldots, p)$ with respect to the new coordinates.

After the above changes of coordinates, the ideal $\mathfrak{j}$ is determined by the equations $\alpha_1^{r_1} = \cdots = \alpha_p^{r_p} = 1$. Setting

$$\mathcal{F} = \{(\mathrm{e}^{2\pi \mathrm{i} s_1 / r_1}, \ldots, \mathrm{e}^{2\pi \mathrm{i} s_p / r_p}, 1, \ldots, 1) \colon \boldsymbol{s} \in \mathbb{N}^p, s_1 < r_1, \ldots, s_p < r_p\},$$

it follows that $\langle M \rangle = \mathcal{F}\ \mathrm{e}^{\mathcal{L}}$. Rewriting $\mathcal{F}$ with respect to the original coordinates now completes the computation of $\langle M \rangle$.

4.3. The algebraic group generated by a single matrix

Let us now consider the case when $\mathcal{M}$ consists of a single arbitrary matrix M. Then we first compute the multiplicative Jordan decomposition of M. Modulo a change of basis of $\mathbb{F}^n$, this means that

$$M = D\,U = U\,D,$$

where $D = M_s$ and $U = M_u$ are the *semi-simple* and *unipotent* parts of M:

$$D = \begin{pmatrix} \lambda_1 I_{n_1} & & \\ & \ddots & \\ & & \lambda_k I_{n_k} \end{pmatrix}, U = \begin{pmatrix} J_{n_1} & & \\ & \ddots & \\ & & J_{n_k} \end{pmatrix},$$

where

$$J_n = \begin{pmatrix} 1 & 1 & & \\ & 1 & \ddots & \\ & & \ddots & 1 \\ & & & 1 \end{pmatrix}.$$

PROPOSITION 28. *We have* $\langle U \rangle = \{\exp(\mu \log U) : \mu \in \mathbb{F}\}$.

REMARK. Notice that $f(N) = \sum_{i=0}^{\infty} f_i\, N^i = \sum_{i=0}^{n-1} f_i\, N^i$ is well-defined for power series $f \in \mathbb{F}[[z]]$ and nilpotent matrices $N \in \mathsf{Mat}_n(\mathbb{F})$; in this case, $f = (1+z)^{\mu}$ and $N = U - 1$.

PROOF. The assertion is clear if $U = I_n$, so assume $U \neq I_n$. Let $\mathcal{X} = \{\exp(\mu \log U) : \mu \in \mathbb{F}\}$. We clearly have $\langle U \rangle \subseteq \mathcal{X}$, since $\mathcal{X}$ is a closed algebraic group which contains U. Moreover, the set $U, U^2, U^3, \ldots$ is infinite, so $\dim \langle U \rangle \geqslant 1$. Since $\mathcal{X}$ is irreducible and $\dim \mathcal{X} = 1$, we conclude that $\langle U \rangle = \mathcal{X}$. □

PROPOSITION 29. *We have* $\langle M \rangle = \langle D \rangle \langle U \rangle$.

PROOF. Since $\langle M \rangle$ is a commutative group, [Hum81, Theorem 15.5] implies that $\langle M \rangle = \langle M \rangle_s \langle M \rangle_u$, where $\langle M \rangle_s = \{N_s : N \in \langle M \rangle\}$ and $\langle M \rangle_u = \{N_u : N \in \langle M \rangle\}$ are closed subgroups of $\langle M \rangle$. Now $\langle D \rangle$ and $\langle U \rangle$ are closed subgroups of $\langle M \rangle_s$ resp. $\langle M \rangle_u$, so $\langle D \rangle \langle U \rangle$ is a closed subgroup of $\langle M \rangle$. Since $M \in \langle D \rangle \langle U \rangle$, it follows that $\langle M \rangle = \langle D \rangle \langle U \rangle$. □

COROLLARY 30. *If* $\langle D \rangle = \mathcal{F}\,\mathrm{e}^{\mathcal{L}}$, *then* $\langle M \rangle = \mathcal{F}\,\mathrm{e}^{\mathcal{L} + \mathbb{F} \log U}$.

4.4. Membership testing for the connected component

In order to compute the closure of the product of a finite number of algebraic groups of the form $\mathcal{F}\,\mathrm{e}^{\mathcal{L}}$, an important subproblem is to test whether a given matrix $M \in \mathsf{GL}_n(\mathbb{F})$ belongs to $\mathrm{e}^{\mathcal{L}}$.

We first observe that $M \in \mathrm{e}^{\mathcal{L}}$ implies $\langle M \rangle \subseteq \mathrm{e}^{\mathcal{L}}$. After the computation of $\mathcal{F}'$ and $\mathcal{L}'$ with $\langle M \rangle = \mathcal{F}'\,\mathrm{e}^{\mathcal{L}'}$ it therefore suffices to check that $\mathcal{L}' \subseteq \mathcal{L}$ and $\mathcal{F}' \subseteq \mathrm{e}^{\mathcal{L}}$. In fact, it suffices to check whether $M' \in \mathrm{e}^{\mathcal{L}}$, where M' is the unique matrix in $\mathcal{F}'$ with $M \in M'\,\mathrm{e}^{\mathcal{L}'}$. Modulo a suitable base change, we have thus reduced the general problem to the case when M is a diagonal matrix whose eigenvalues are all roots of unity.

Assume that $M \in \mathrm{e}^{\mathcal{L}}$ and $\ell \in \mathcal{L}$ are such that $M \in \mathrm{e}^{\mathbb{C}\ell}$. Since M and ℓ commute, it follows that M and ℓ can be diagonalized w.r.t. a common basis. The elements of this basis are elements of the different eigenspaces of M. In other words, if $M = \mathrm{Diag}(\lambda_1 I_{n_1}, \ldots, \lambda_k I_{n_k})$ with pairwise distinct λ_i, then $P^{-1}\,\ell\,P$ is diagonal for some block matrix $P = \mathrm{Diag}(P_1, \ldots, P_k)$ with $P_i \in \mathsf{GL}_{n_i}(\mathbb{F})$ for each i. It follows that $\ell = \mathrm{Diag}(\ell_1, \ldots, \ell_k)$ for certain $\ell_i \in \mathsf{Mat}_{n_i}(\mathbb{F})$. Without loss of generality, we may therefore replace $\mathcal{L}$ by the intersection of $\mathcal{L}$ with $\mathrm{Diag}(\mathsf{Mat}_{n_1}(\mathbb{F}), \ldots, \mathsf{Mat}_{n_k}(\mathbb{F}))$.

From now on, we assume that the above two reductions have been made. Let $\ell = \mathrm{Diag}(\mu_1, \ldots, \mu_n)$ be a diagonal matrix in $\mathcal{L}$. By lemma 27, we have $M \in \mathrm{e}^{\mathbb{C}\ell}$ if and only if any $\mathbb{Z}$-linear relation $\boldsymbol{l} \cdot \boldsymbol{\mu} = 0$ induces a relation $\boldsymbol{\lambda}^{\pi(\boldsymbol{l})} = 1$, where $\pi(\boldsymbol{l}) = (l_1 + \cdots + l_{n_1}, \ldots, l_{n-n_k} + \cdots + l_n)$. Now consider a random matrix R in $\mathcal{L}$, i.e. a linear combination of the basis elements with small random integer coefficients. We compute its blockwise Jordan normal form $J = P^{-1} R P$ so that $P \in \mathrm{Diag}(\mathsf{GL}_{n_1}(\mathbb{F}), \ldots, \mathsf{GL}_{n_k}(\mathbb{F}))$ and let ℓ be the restriction of J to the diagonal. We have $M \in \mathrm{e}^{\mathbb{C}\ell} \Leftrightarrow M \in \mathrm{e}^{\mathbb{C}J} \Leftrightarrow M = P M P^{-1} \in \mathrm{e}^{\mathbb{C}R}$. Computing a basis for the $\mathbb{Z}$-linear relations of the form $\boldsymbol{l} \cdot \boldsymbol{\mu} = 0$ using **EH3**, the above criterion now enables us to check whether $M \in \mathrm{e}^{\mathbb{C}R}$.

If the check whether $M \in \mathrm{e}^{\mathbb{C}R}$ succeeds, then we are clearly done. Otherwise, since R was chosen in a random way, the relation $\boldsymbol{l} \cdot \boldsymbol{\mu}$ is very likely to be satisfied for all possible choices of $R \in \mathcal{L}$ (up to permutations of coordinates inside each block). Indeed, the R for which this is not the case lie on a countable union $\mathcal{U}$ of algebraic variety of a lower dimension, so $\mathcal{U}$ has measure 0. Heuristically speaking, we may therefore conclude that $M \notin \mathrm{e}^{\mathcal{L}}$ if the check fails (at least temporarily, modulo some final checks when the overall computation of $\langle \mathcal{M} \rangle$ will be completed).

Theoretically speaking, we may perform the above computations with $R' = \sum_{B \in \mathcal{B}} \alpha_B B$ instead of R, where $\mathcal{B}$ is a basis of $\mathcal{L}$ and the α_B are formal parameters. We then check whether the relation $\boldsymbol{l} \cdot \boldsymbol{\mu}'$ is still satisfied for the analogue $\ell' = \mathrm{Diag}(\mu_1', \ldots, \mu_n')$ of ℓ. If so, then we are sure that $M \notin \mathrm{e}^{\mathcal{L}}$. Otherwise, we keep trying with other random elements of $\mathcal{L}$.

It is likely that a more efficient theoretical algorithm can be designed for testing $\mathbb{Z}$-linear relations between the eigenvalues of elements in $\mathcal{L}$. One of the referees suggested to use similar methods as in [Mas88, Ber95, CS98]. However, we did not study this topic in more detail, since our final algorithm for the computation of Galois groups will be based on heuristics anyway. We also notice that a "really good" random number generator should actually *never* generate points which satisfy non-trivial algebraic relations.

4.5. Computing the closure of $\mathcal{M}$

A Lie algebra $\mathcal{L}$ is said to be *algebraic*, if it is the Lie algebra of some algebraic group, i.e. if $\mathrm{e}^{\mathcal{L}}$ is an algebraic subset of $\mathsf{GL}_n(\mathbb{F})$. It is classical [Bor91, Corollary 7.7] that the smallest Lie algebra generated by a finite number of algebraic Lie algebras is again algebraic. The Lie algebras we will consider in our algorithms will always assumed to be algebraic. Given a finite number $\mathcal{L}_1, \ldots, \mathcal{L}_l$ of algebraic Lie algebras and a basis $\mathcal{B}$ for $\mathcal{L}_1 + \cdots + \mathcal{L}_l$, it is easy to enrich $\mathcal{B}$ so that $\mathcal{L} = \mathsf{Vect}(\mathcal{B})$ is a Lie algebra: as long as $[\boldsymbol{b}_1, \boldsymbol{b}_2] \notin \mathcal{L}$ for two elements $\boldsymbol{b}_1, \boldsymbol{b}_2 \in \mathcal{B}$, we add $[\boldsymbol{b}_1, \boldsymbol{b}_2]$ to $\mathcal{B}$. By what precedes, the computed Lie algebra $\mathcal{L}$ is again algebraic.

Putting together the ingredients from the previous sections, we now have the following algorithm for computing the smallest closed algebraic group $\langle \mathcal{M} \rangle$ which contains $\mathcal{M}$.

Algorithm `Closure`$(\mathcal{M})$
Input: A subset $\mathcal{M} = \{M_1, \ldots, M_m\}$ of $\mathsf{GL}_n(\mathbb{F})$
Output: a numeric approximation of $\langle \mathcal{M} \rangle$

Step 1. [Initialize algorithm]
Compute $\langle M_i \rangle = \mathcal{F}_i \, \mathrm{e}^{\mathcal{L}_i}$ for each $i \in \{1, \ldots, m\}$
Let $\mathcal{F} := \mathcal{F}_1 \cup \cdots \cup \mathcal{F}_m$ (notice that $1 \in \mathcal{F}$)
Let $\mathcal{L} := \mathsf{Lie}(\mathcal{L}_1 + \cdots + \mathcal{L}_m)$

Step 2. [Closure]
While there exists an $N \in \mathcal{F} \setminus \{1\}$ with $N \mathcal{L} N^{-1} \nsubseteq \mathcal{L}$ set $\mathcal{L} := \mathsf{Lie}(\mathcal{L} + N \mathcal{L} N^{-1})$
While there exists an $N \in \mathcal{F} \setminus \{1\}$ with $N \in \mathrm{e}^{\mathcal{L}}$ set $\mathcal{F} := \mathcal{F} \setminus \{N\}$
While there exists $N \in \mathcal{F}^2$ with $N \notin \mathcal{F}\, \mathrm{e}^{\mathcal{L}}$ do
Compute $\langle N \rangle = \mathcal{F}' \mathrm{e}^{\mathcal{L}'}$
If $\mathcal{L}' \nsubseteq \mathcal{L}$ then set $\mathcal{L} := \mathsf{Lie}(\mathcal{L} + \mathcal{L}')$, quit loop and repeat step 2
Otherwise, set $\mathcal{F} := \mathcal{F} \cup \{N\}$
Return $\mathcal{F}\, \mathrm{e}^{\mathcal{L}}$

The termination of this algorithm relies on a lemma, whose proof was kindly communicated to the author by J.-Y. Hée.

LEMMA 31. *Let $\mathcal{H}$ be a closed algebraic subgroup of $\mathsf{GL}_n(\mathbb{C})$ and let $M_1, \ldots, M_m \in \mathsf{GL}_n(\mathbb{C})$ be a finite number of matrices in the normalizer of $\mathcal{H}$. Denote by $\mathcal{G}$ the group generated by $\mathcal{H}$ and $M_1, \ldots, M_m$. If all elements in $\mathcal{G}/\mathcal{H}$ have finite order, then $\mathcal{G}/\mathcal{H}$ is finite.*

PROOF. In the case when $\mathcal{H} = \{1\}$, the result is classical [Dix71, Theorem 9.2]. In the general case, the normalizer $\mathcal{N}$ of $\mathcal{H}$ is a closed algebraic subgroup of $\mathsf{GL}_n(\mathbb{C})$ and $\mathcal{H}$ is a normal subgroup of $\mathcal{N}$. By [Bor91, Theorem 6.8 and Proposition 1.10], it follows that $\mathcal{N}/\mathcal{G}$ is an affine algebraic group which is isomorphic to a closed algebraic matrix group. This reduces the general case to the special case when $\mathcal{H} = \{1\}$. □

THEOREM 32. *There exists an $\varepsilon_0 \in \mathbb{N}^{>} 2^{\mathbb{Z}}$ such that, for every $\varepsilon \in \mathbb{N}^{>} 2^{\mathbb{Z}}$ with $\varepsilon < \varepsilon_0$, the set $\mathcal{F}\, \mathrm{e}^{\mathcal{L}}$ returned by* `Closure`*, considered as a subset of $\mathsf{GL}_n(\mathbb{C}^{\mathrm{eff}})$, coincides with the smallest closed algebraic subgroup $\langle \mathcal{M} \rangle$ of $\mathsf{GL}_n(\mathbb{C}^{\mathrm{eff}})$ which contains $\mathcal{M}$.*

PROOF. Clearly, the dimension of $\mathcal{L}$ increases throughout the execution of the algorithm, so it remains ultimately constant. At this point, the set $\mathcal{F}$ will keep growing and the lemma implies that $\mathcal{F}$ ultimately stabilizes. When this happens, $\mathcal{F}$ is closed under multiplication modulo $\mathrm{e}^{\mathcal{L}}$, as well as under multiplicative inverses, since each element in $\mathcal{F}$ has finite order modulo $\mathrm{e}^{\mathcal{L}}$. We conclude that $\mathcal{F}\, \mathrm{e}^{\mathcal{L}}$ is indeed the smallest closed algebraic subgroup of $\mathsf{GL}_n(\mathbb{F})$ which contains $\mathcal{M}$, provided that the approximate zero-test always returns the right result.

In order to prove the correctness at a sufficient precision, we assume that we use the theoretic membership test from section 4.4 and that the random number generator successively generates the same random numbers each time we relaunch the algorithm at a higher precision. Now consider the trace of the execution of our algorithm when using an infinite precision. Let ε_0 be a sufficient precision such that all zero-tests in this execution tree are still correct when we replace the infinite precision by a precision $\varepsilon < \varepsilon_0$. Then the trace of the execution any finite precision $\varepsilon < \varepsilon_0$ coincides with the trace of the execution at infinite precision. This completes the proof. □

REMARK 33. The main improvement of the algorithm `Closure` w.r.t. the algorithm from [DJK03] lies in the more efficient treatment of the connected component (using linear algebra). On the other hand, the mere enumeration of representatives in each connected component can be very unefficient (although a Gröbner basis might be of the same size). Fortunately, we will see in the next sections how to remove this drawback.

Assume now that $\mathcal{M}$ is the set of generators for $\mathcal{G}_{L,\boldsymbol{h}^{z_0}}$ as computed in theorem 20. Assume that we have computed a reasonable candidate $\mathcal{F}\, \mathrm{e}^{\mathcal{L}}$ for $\langle \mathcal{M} \rangle$, expressed in the original basis corresponding to $\boldsymbol{h}^{z_0}$. We still have to reconstruct $\tilde{\mathcal{F}} \subseteq \mathsf{GL}_n(\mathbb{K})$ and $\tilde{\mathcal{L}} = \mathsf{Vect}(\tilde{\mathcal{B}})$ with $\tilde{\mathcal{B}} \in \mathsf{Mat}_n(\mathbb{K})$ such that $\mathcal{F}\, \mathrm{e}^{\mathcal{L}} \cap \mathsf{GL}_n(\mathbb{K}) = \tilde{\mathcal{F}}\, \mathrm{e}^{\tilde{\mathcal{L}}} \cap \mathsf{GL}_n(\mathbb{K})$.

In the case of $\tilde{\mathcal{L}}$, by selecting a suitable basis of $\mathsf{Mat}_n(\mathbb{F})$, we may consider $\tilde{\mathcal{B}}$ as a big $d \times n^2$ matrix whose first d columns are linearly independent. We compute the row-echelon form of this basis:

$$E = \begin{pmatrix} 1 & & & * & \cdots & * \\ & \ddots & & \vdots & & \vdots \\ & & 1 & * & \cdots & * \end{pmatrix}.$$

The entries of E must be in $\mathbb{K}$: provided that $\tilde{\mathcal{L}}$ is indeed generated by a basis of matrices with entries in $\mathbb{K}$, the row-echelon form of this second basis coincides with E. It therefore suffices to reconstruct the entries of E using the LLL-algorithm.

In the case of a matrix $M \in \mathcal{F}$, the set $M\,\mathrm{e}^{\tilde{\mathcal{L}}}$ is an algebraic variety of dimension d over $\mathbb{K}$. Now choose $\tilde{M} \in M\,\mathrm{e}^{\tilde{\mathcal{L}}}$ close to M in such a way that d independent coordinates of $\tilde{M}$ are all in $\mathbb{Q} \subseteq \mathbb{K}$. Then the other coordinates of $\tilde{M}$, considered as elements of $\mathbb{C}^{\mathrm{eff}}$, are easily found using Newton's method. Since $M\,\mathrm{e}^{\tilde{\mathcal{L}}}$ is an algebraic variety, these other coordinates are actually in $\mathbb{K}$, and we reconstruct them using the LLL-algorithm.

4.6. Fast computations with the connected components

The algorithm `Closure` from the previous section is quite inefficient when the set $\mathcal{F}$ becomes large. It is therefore useful to seek for a better computational representation of $\mathcal{F}$. For finite groups $\mathcal{G}$, one classical idea is to search for a sequence of subgroups

$$1 = \mathcal{G}_0 \subsetneq \mathcal{G}_1 \subsetneq \cdots \subsetneq \mathcal{G}_k = \mathcal{G} \tag{9}$$

such that the indices $\mathcal{G}_i{:}\mathcal{G}_{i-1}$ are small. Then we may represent elements in $\mathcal{F}$ by sequences $(a_1, \ldots, a_k)$ with $a_i \in \mathcal{G}_i/\mathcal{G}_{i-1}$ for each i. This representation is particularly useful if $\mathcal{F}$ operates on a set $\mathcal{S}$ and if there exists points $a_1, \ldots, a_k$ in $\mathcal{S}$ such that

$$\mathcal{G}_i = S_{a_1, \ldots, a_{k-i}}$$

is the stabilizer of the set $\{a_1, \ldots, a_{k-i}\}$ for each i. Then the set $S_{a_1,\ldots,a_{i-1}}/S_{a_1,\ldots,a_i}$ corresponds to the orbit of a_i while leaving $a_1, \ldots, a_{i-1}$ fixed [Sim70, Sim71].

In the case of matrix groups, one often takes $\mathbb{F}^n$ for $\mathcal{S}$ [MO95]. However, this approach only yields interesting results when there exist non-trivial invariant subspaces under the action of the group, which will usually not be the case for us (otherwise we may factor L and consider smaller problems). A theoretical way out of this is to also consider the action of $\mathcal{F}$ on exterior powers $\wedge^p\ \mathbb{F}^n$. However, this approach is very expensive from a computational point of view. In our more specific context of matrices with complex entries, we will therefore combine two other approaches: non-commutative lattice reduction and the operation of $\mathcal{F}$ on $\mathsf{Mat}_n(\mathbb{F})/\mathrm{e}^{\mathcal{L}}$ via conjugations $M \mapsto M N M^{-1}$.

The algebra $\mathsf{Mat}_n(\mathbb{F})$ admits a natural (multiplicative) norm, given by

$$\|M\| = \sup\,\{|M V|: V \in \mathbb{F}^n, |V| = 1\},$$

where $|\cdot|$ stands for the Euclidean norm on $\mathbb{F}^n$. If $\mathcal{G} = \langle \mathcal{M} \rangle / \mathrm{e}^{\mathcal{L}}$ is finite, this enables us to construct $\mathcal{G}_0 = 1, \mathcal{G}_1, \ldots, \mathcal{G}_k$ as in (9) as follows. Assuming that $\mathcal{G}_0, \ldots, \mathcal{G}_{i-1}$ have been constructed, we consider a matrix $M_i \in \langle \mathcal{M} \rangle \setminus \mathcal{G}_{i-1}\,\mathrm{e}^{\mathcal{L}}$ for which $\|M_i - 1\|$ is minimal, and let $\mathcal{G}_i$ be the set generated by $M_i\,\mathrm{e}^{\mathcal{L}}$ and $\mathcal{G}_{i-1}$ in $\mathcal{G}$. This construction allows us to rapidly identify a big commutative part of $\mathcal{G}$. More precisely, we have

Proposition 34. *Let $A, B \in \mathsf{GL}_n(\mathbb{F})$ be such that $\varepsilon = \|A - 1\| < 1$ and $\delta = \|B - 1\| < 1$. Then we have*

$$\|A B A^{-1} B^{-1} - 1\| \leqslant B(\delta, \varepsilon) = \frac{2\,\varepsilon^2}{1 - \varepsilon} + \frac{2\,\delta^2}{1 - \delta} + \frac{4\,\varepsilon\,\delta}{(1 - \varepsilon)\,(1 - \delta)}.$$

PROOF. Writing $A = 1 + \Delta$ and $B = 1 + \mathrm{E}$, we expand $A^{-1} = 1 - \Delta + \Delta^2 + \cdots$ and $B^{-1} = 1 - \mathrm{E} + \mathrm{E}^2 + \cdots$ in $A\,B\,A^{-1}\,B^{-1}$. This yields a non-commutative power series in Δ and E whose terms in $1, \Delta$ and E vanish. It follows that

$$\|A\,B\,A^{-1}\,B^{-1} - 1\| \leqslant (1+\delta)\,(1+\varepsilon)\,\frac{1}{1-\delta}\,\frac{1}{1-\varepsilon} - 1 - 2\,\delta - 2\,\varepsilon = B(\delta,\varepsilon). \qquad \square$$

The proposition implies that $\|A\,B\,A^{-1}\,B^{-1} - 1\| < \min(\varepsilon, \delta)$ whenever $\max(\varepsilon, \delta) < 5 - \sqrt{24}$. Now take $A = M_i$ and $B = M_j$ with $i < j$, where the M_i are as above. Then it follows that A and B commute whenever $B(\delta, \varepsilon) < \|M_1 - 1\|$. What is more, proposition 34 shows that taking commutators is a convenient way to construct matrices close to identity from a set of non-commutative generators.

However, from the effective point of view, we will not compute the *exact* computation of minimal representatives M_i in cosets of algebraic groups in detail. We will rather simulate such a computation in a way which is sufficient for our purpose. If $M \in \mathcal{F}$ is such that $\|M - 1\|$ is small, then we will also try to use the fact that the centralizer C_M of M is often a big subgroup of $\langle \mathcal{M} \rangle / \mathrm{e}^{\mathcal{L}}$, so the orbit of $N \mapsto N^{-1}\,M\,N$ is small.

4.7. Non-commutative lattice reduction

Let $\mathcal{G}$ be a closed algebraic subgroup of $\mathsf{Mat}_n(\mathbb{F})$ with associated Lie-algebra $\mathcal{L}$. In this section, we will show how to compute efficiently with elements of the finite group $\mathcal{H} = \mathcal{G}/\mathrm{e}^{\mathcal{L}}$. Until the very end of this section, we assume that $\mathcal{G}$ is included in the connected component of the normalizer of $\mathrm{e}^{\mathcal{L}}$ in $\mathsf{Mat}_n(\mathbb{F})$. We denote by $\mathcal{N}$ the Lie algebra of this connected component. By [Hum81, Theorem 13.3], we have $\mathcal{N} = \{N \in \mathsf{Mat}_n(\mathbb{F}) : [N, \mathcal{L}] \subseteq \mathcal{L}\}$.

Orthogonal projection. Let $M \in \mathcal{G}$ and recall that M belongs to the normalizer of $\mathrm{e}^{\mathcal{L}}$. If $\|M - 1\| < 1$, then $X = \log(1 + (M - 1))$ also belongs to the normalizer of $\mathrm{e}^{\mathcal{L}}$. Since $M \in \mathrm{e}^{\mathbb{F}X}$ lies in the connected component of this normalizer, we have $X \in \mathcal{N}$. Now consider the orthogonal supplement $\mathcal{L}^{\perp}$ of $\mathcal{L}$ for the Hermitian product on $\mathsf{Mat}_n(\mathbb{F})$. We define $\pi_{\mathcal{L}}(M) = \mathrm{e}^{Y}$, where Y is the orthogonal projection of X on $\mathcal{L}^{\perp}$. From $[X, \mathcal{L}] \subseteq \mathcal{L}$, it follows that $\pi_{\mathcal{L}}(M) \in M\,\mathrm{e}^{\mathcal{L}}$, and we denote $\|M\|_{\mathcal{L}} = \|\pi_{\mathcal{L}}(M)\|$. Since $\mathrm{e}^{\mathcal{N}}$ is connected, the function $M \mapsto \log M$ may actually be analytically continued to a multivalued function $\mathrm{e}^{\mathcal{N}} \to \mathcal{N}$. After choosing branch cuts (the way this is done is not crucial for what follows, provided that we make the standard choice for M with $\|M - 1\| < 1$), this allows us to extend the definitions of $\pi_{\mathcal{L}}$ and $\|\cdot\|_{\mathcal{L}}$ to the case when $\|M - 1\| \geqslant 1$.

Representation of the elements in $\mathcal{H}$. Let $X \in \mathcal{N}$ and $M = \mathrm{e}^{X}$ be such that

- $M\,\mathrm{e}^{\mathcal{L}} \in \mathcal{H}$.
- $M\,\mathrm{e}^{\mathcal{L}}$ generates $(\mathrm{e}^{\mathbb{F}X} \cap \mathcal{G})/\mathrm{e}^{\mathcal{L}}$.
- $\|M\|_{\mathcal{L}} \leqslant \|M^k\|_{\mathcal{L}}$ whenever $M^k \in M^{\mathbb{Z}}$ is another such generator.

Let $p_1 \cdots p_l$ be the prime-decomposition of the order of M modulo $\mathrm{e}^{\mathcal{L}}$, with $p_1 \geqslant \cdots \geqslant p_l$. Let $A_0 = X$ and $A_i = M^{p_1 \cdots p_i}$ for all $i \in \{1, \ldots, l\}$. Let $\mathcal{H}_i$ be the subgroup of $\mathcal{H}$ of elements which commute with A_i modulo $\mathrm{e}^{\mathcal{L}}$, so that $\mathcal{H}_0 \subseteq \cdots \subseteq \mathcal{H}_l = \mathcal{H}$. For $i \in \{1, \ldots, r\}$, we represent elements in the quotient $\mathcal{H}_i/\mathcal{H}_{i-1}$ by elements in the orbit of the action $\Phi_{A_i}: N \mapsto A_i\,N\,A_i^{-1}$ modulo $\mathcal{H}_{i-1}$. Since $[X, \mathcal{L}] \subseteq \mathcal{L}$, the set $\mathcal{L}' = \mathcal{L} \oplus \mathbb{F}\,X$ is a Lie algebra whose normalizer contains $\mathcal{H}_0$. Consequently, $\mathcal{H}_0 \cong M^{\mathbb{Z}} \times \mathcal{H}_0/\mathrm{e}^{\mathcal{L}'}$, and we represent elements in $\mathcal{H}_0$ by products $M^k\,P$, with $k \in \mathbb{Z}$ and $P\,\mathrm{e}^{\mathcal{L}'} \in \mathcal{H}_0/\mathrm{e}^{\mathcal{L}'}$. The elements in $\mathcal{H}_0/\mathrm{e}^{\mathcal{L}'}$ are represented in a similar way as the elements in $\mathcal{H}$, using recursion. The successive matrices M for $\mathcal{G}/\mathrm{e}^{\mathcal{L}}$, $\mathcal{H}_0/\mathrm{e}^{\mathcal{L}'}$, etc. will be called a *basis* for $\mathcal{H}$. A basis $(B_1, \ldots, B_m)$ is said to be *sorted* if $\|B_1\|_{\mathcal{L}} \leqslant \cdots \leqslant \|B_m\|_{\mathcal{L}}$.

Adding new elements to a basis. Let $(B_1, \ldots, B_m)$ be a sorted basis for $\mathcal{H} = \mathcal{G}/\mathrm{e}^{\mathcal{L}}$ and assume that we want to compute the extension $\hat{\mathcal{G}} = \langle \mathcal{G}, N \rangle$ of $\mathcal{G}$ by a new matrix N. Whenever we hit an element $\hat{M}\,\mathrm{e}^{\mathcal{L}} \in \hat{\mathcal{H}} = \hat{\mathcal{G}}/\mathrm{e}^{\mathcal{L}}$ with $\|\hat{M}\|_{\mathcal{L}} < \|B_1\|_{\mathcal{L}}$ during our computations, then we start the process of basis reduction, which is described below. Whenever we find an element in $\hat{\mathcal{L}} \setminus \mathcal{L}$, then we abort all computations and return this element (indeed, in that case, we may continue with the closure of the connected component in `Closure`).

Let $M = B_1$, X, etc. be as above. We start by computing the orbit of $\hat{\mathcal{G}}$ modulo $\mathcal{H}_{r-1}$ for Φ_{A_r}. Whenever we hit an element $P \neq 1$ (modulo $\mathrm{e}^{\mathcal{L}}$) with $\|B_1 P B_1^{-1} P^{-1}\|_{\mathcal{L}} < \|B_1\|_{\mathcal{L}}$ or $\|P\|_{\mathcal{L}} < \|B_1\|_{\mathcal{L}}$, then we start the process of basis reduction. Otherwise, we obtain a finite orbit, together with a finite number of matrices by which we have to extend $\mathcal{H}_{r-1}$. We keep doing this using the same method for $\mathcal{H}_{r-1}$ until $\mathcal{H}_1$.

At the end, we still have to show how to extend $\mathcal{H}_0$ with a new matrix $\tilde{N}$. Now recursive application of the algorithm to $\tilde{\mathcal{H}} = \mathcal{H}_0/\mathrm{e}^{\mathcal{L}'}$ and $\tilde{N}$ yields a sorted basis $\tilde{B}_1, \ldots, \tilde{B}_{\tilde{m}}$. When keeping track of the corresponding powers of e^X during the computations, we also obtain a finite system of generators for $\hat{\mathcal{G}} \cap \mathrm{e}^{\mathbb{F}X}$. Using g.c.d. computations we either obtain a minimal generator $\hat{B}_1$ or a new element in the connected component. In the first case, we return $(\hat{B}_1, \tilde{B}_2, \ldots, \tilde{B}_{\tilde{m}})$ if $\|\hat{B}_1\|_{\mathcal{L}} < \|\tilde{B}_2\|_{\mathcal{L}}$ and apply basis reduction otherwise.

Basis reduction. Let $B_1, \ldots, B_m$ be in $\mathcal{G}$ with $\|B_1\|_{\mathcal{L}} \leqslant \cdots \leqslant \|B_m\|_{\mathcal{L}}$. We call $(B_1, \ldots, B_m)$ a *raw basis*. In the above algorithms, raw bases occur when we are given an ordered basis $(B_2, \ldots, B_m)$ for $\mathcal{G}$, and we find a new element B_1 with $\|B_1\|_{\mathcal{L}} < \|B_2\|_{\mathcal{L}}$.

Using the above base extension procedure, we may transform a raw basis $(B_1, \ldots, B_m)$ into a basis for $\mathcal{G}$: starting with (B_1), we successively add $B_2, \ldots, B_m$. However, it is more efficient to reduce $(B_1, \ldots, B_m)$ first. More precisely, let us now describe a procedure which tries to replace $(B_1, \ldots, B_m)$ by a better raw basis $(\tilde{B}_1, \ldots, \tilde{B}_{\tilde{m}})$, with $\langle \tilde{B}_1, \ldots, \tilde{B}_{\tilde{m}} \rangle = \langle B_1, \ldots, B_m \rangle$, and whose elements are closer to identity. Of course, we may always return the original basis if a better one could not be found.

We first test whether all basis elements are roots of unity modulo $\mathcal{L}$. If not, then we found a new element in the connected component. We next test whether there exist i, j with $\|B_i B_j B_i^{-1} B_j^{-1}\|_{\mathcal{L}} < \|B_1\|_{\mathcal{L}}$, in which case we keep adding the smallest such commutator to the basis. Whenever this stops, we write $B_1 = \mathrm{e}^{X_1}, \ldots, B_m = \mathrm{e}^{X_m}$ with $X_1, \ldots, X_m \in \mathcal{L}^{\perp}$ and consider all lattice reductions $X_i \leftarrow X_i + k\, X_j$ $(k \in \mathbb{Z})$ proposed by the LLL-algorithm in the commutative vector space $\mathcal{L}^{\perp}$. Whenever $0 < \|B_i B_j^k\|_{\mathcal{L}} < \|B_i\|$, for one such reduction, then we perform the corresponding reduction $B_i \leftarrow B_i B_j^k$ on our basis and keep repeating the basis reduction process.

The general case. We still have to show how to deal with the case when $\mathcal{G}$ is not included in the connected component $\mathrm{e}^{\mathcal{N}}$ of the normalizer of $\mathrm{e}^{\mathcal{L}}$ in $\mathsf{Mat}_n(\mathbb{F})$. In that case, we start with the computation of a basis for $\mathcal{N}$, using linear algebra. Since $\mathrm{e}^{\mathcal{N}} \cap \mathcal{G}$ is a normal subgroup of $\mathcal{G}$, we have $\mathcal{G} \cong (\mathcal{G}/\mathrm{e}^{\mathcal{N}})\,(\mathrm{e}^{\mathcal{N}} \cap \mathcal{G})$. Now we have explained above how to compute with elements in $\mathrm{e}^{\mathcal{N}} \cap \mathcal{G}$. If $\mathcal{N} \supsetneq \mathcal{L}$, then may use recursion for computations in the finite group $\mathcal{G}/\mathrm{e}^{\mathcal{N}}$. If $\mathcal{N} = \mathcal{L}$, then elements in $\mathcal{G}/\mathrm{e}^{\mathcal{N}}$ have necessarily small order, so we simply list the elements of $\mathcal{G}/\mathrm{e}^{\mathcal{L}}$.

5. Conclusion and final notes

We hope that we provided convincing evidence that analytic methods may be used for the efficient computation of differential Galois groups and related problems like the factorization of linear differential operators.

The two main remaining challenges are the concrete implementation of the algorithms presented here (as part of a more general library for the computation with analytic functions such as [vdHea05]) and the development of *a priori* or *a posteriori* methods for ensuring the correctness of the computed result. Some ideas into this direction are as follows:

- Use theoretical bounds on the number of connected components of the computed Galois group and related bounds on the sizes of the basis elements in **EH2** and **EH3**. See [DJK03, Section 3.2] for some results.

- Use the classification theory for algebraic groups in order to gather more information about the computed Galois group $\mathcal{G}$. In particular, it is useful to compute the radical (or unipotent radical) of $\mathcal{G}$, thereby reducing the study of $\mathcal{G}$ to the study of a finite group, a semisimple (or reductive) group and a solvable (or unipotent) group [Hum81, page 125]. We refer to [dG00] for computational aspects of the corresponding Lie algebras.

- Use the classical theory of invariant subspaces in symmetric products or exterior powers as an *a posteriori* correctness check and search for an effective version of Chevalley's theorem [Hum81, Theorem 11.2]. One may start with generalizing [vHW97, CS98] and notice that a better knowledge of the Galois group $\mathcal{G}$ helps to further restrict the number of monomials (i.e. "generalized exponents") to be considered. Indeed, if $\mathcal{H}$ is an arbitrary algebraic subgroup of $\mathcal{G}$, for which the ring of invariants is easy to compute, then the invariants for $\mathcal{G}$ must be searched in this ring. Also, their are known algorithms for computing the invariants for certain types of algebraic groups, like linearly reductive groups [Der99].

- The representation for algebraic groups $\mathcal{G}$ we used in section 4 is efficient for computations (we merely do linear algebra in dimension n^2, lattice reduction and computations with small finite groups). Nevertheless, it may be interesting to reconstruct the algebraic equations for $\mathcal{G}$ and search for equations which are particularly sparse with respect to suitably chosen coordinates. For instance, a big cyclic group admits a particularly nice (resp. large) Gröbner basis w.r.t. well chosen (resp. badly chosen) coordinates. Conversely, it may be interesting to switch back from a Gröbner basis representation to our representation.

- Carefully identify those parts of the algorithm which either prove or disprove certain matrices to belong to the Galois group. For instance, we know that all Stokes matrices are unipotent. Given a non-zero transcendental number λ, we may then reliably conclude that a Stokes matrix of the form $\begin{pmatrix} 1 & \lambda \\ 0 & 1 \end{pmatrix}$ generates the group $\{\begin{pmatrix} 1 & \alpha \\ 0 & 1 \end{pmatrix} : \alpha \in \mathbb{K}\}$.

- An interesting idea to get rid of the transcendental part of the computations might be to quotient the values of the functions in our basis $\boldsymbol{h}$ of solutions by the action of the Galois group. For instance, if z_0 and z_1 are close regular points in $\mathbb{K}$, is it true that the orbit of $\boldsymbol{h}^{z_0}(z_1)$ under the action of the Galois group necessarily contains a point in $\mathbb{K}^n$? This is clearly the case for finite Galois groups and the full Galois group, as well as for the equations $f' = f$ and $(z f')' = 0$. More generally, as soon as $\boldsymbol{h}^{z_0}(z_1)$ becomes more transcendental, its orbit under the action of the Galois group becomes larger, so the likelihood of finding a point in the intersection with $\mathbb{K}^n$ increases.

Besides the above ideas for improving the algorithms, this paper also raises a few other interesting questions:

- Are there more efficient approaches for the reconstruction of elements in $\mathbb{K}$ in section 3.4, both in the cases when $\mathbb{K} = \mathbb{Q}^{\mathrm{alg}}$ and when $\mathbb{K}$ is more general? Also, as pointed out above, we may want to reconstruct equations for $\mathcal{G}$ from the variety.
- Does there exists an efficient membership test in section 4.4 which does not rely on probabilistic arguments?
- Can the approach of section 4 be adapted to the computation of a "basis" for the usual topological closure of a finitely generated matrix group?

Of course, a better mastering of the algorithms in this paper may also lead to more efficient algorithms for other computations which rely on differential Galois theory, like the computation of Liouvillian or other forms of solutions. More generally, our algorithms may be used for other computations with algebraic matrix groups over $\mathbb{C}$ and other fields of characteristic 0. We also expect all results to generalize to holonomic systems of linear partial differential equations.

Acknowledgment. The author would like to thank J.-Y. Hée, W. de Graaf, F. Zara, J. Écalle and the referees for several helpful comments and references.

Bibliography

[BB85] D. Bertrand and F. Beukers. équations différentielles linéaires et majorations de multiplicités. *Ann. Sci. de l'École Norm. Sup.*, 28(4-5):473–494, 1985.

[Bek94] E. Beke. Die Irreduzibilität der homogenen linearen Differentialgleichungen. *Math. Ann.*, 45, 1894.

[Ber95] D. Bertrand. Minimal heights and polarizations on group varieties. *Duke Math. J.*, 80(1):223–250, 1995.

[Bor91] A. Borel. *Linear algebraic groups*. Springer-Verlag, 2nd edition, 1991.

[Bra91] B.L.J. Braaksma. Multisummability and Stokes multipliers of linear meromorphic differential equations. *J. Diff. Eq*, 92:45–75, 1991.

[Bra92] B.L.J. Braaksma. Multisummability of formal power series solutions to nonlinear meromorphic differential equations. *Ann. Inst. Fourier de Grenoble*, 42:517–540, 1992.

[CC90] D.V. Chudnovsky and G.V. Chudnovsky. Computer algebra in the service of mathematical physics and number theory (computers in mathematics, stanford, ca, 1986). In *Lect. Notes in Pure and Applied Math.*, volume 125, pages 109–232, New-York, 1990. Dekker.

[Clu04] T. Cluzeau. *Algorithmique modulaire des équations différentielles linéaires*. PhD thesis, Université de Limoges, 2004.

[CNP93] B. Candelberger, J.C. Nosmas, and F. Pham. *Approche de la résurgence*. Actualités mathématiques. Hermann, 1993.

[CS98] É. Compoint and M. Singer. Computing Galois groups of completely reducible differential equations. *JSC*, 11:1–22, 1998.

[Der99] H. Derksen. Computation of reductive group invariants. *Adv. in Math.*, 141:366–384, 1999.

[dG00] W. de Graaf. *Lie Algebras: Theory and Algorithms*, volume 56 of *North Holland Mathematical Library*. Elsevier science, 2000.

[Dix71] J.D. Dixon. *The structure of linear groups*. Van Nostrand Reinhold Company, 1971.

[DJK03] H. Derksen, E. Jaendel, and P. Koiran. Quantum automata and algebraic groups. Technical Report 2003-39, LIP, École Norm. Sup. de Lyon, 2003. Presented at MEGA 2003; to appear in JSC.

[É85] J. Écalle. *Les fonctions résurgentes I–III*. Publ. Math. d'Orsay 1981 and 1985, 1985.

[É87] J. Écalle. L'accélération des fonctions résurgentes (survol). Unpublished manuscript, 1987.

[É92] J. Écalle. *Introduction aux fonctions analysables et preuve constructive de la conjecture de Dulac*. Hermann, collection: Actualités mathématiques, 1992.

[É93] J. Écalle. Six lectures on transseries, analysable functions and the constructive proof of Dulac's conjecture. In D. Schlomiuk, editor, *Bifurcations and periodic orbits of vector fields*, pages 75–184. Kluwer, 1993.

[Fab85] E. Fabry. *Sur les intégrales des équations différentielles linéaires à coefficients rationnels*. PhD thesis, Paris, 1885.

[Hum81] J.E. Humphreys. *Linear algebraic groups*. Graduate Texts in Math. Springer, 1981.

[Kap57] I. Kaplansky. *An introduction to differential algebra*. Hermann, 1957.

[Kol73] E.R. Kolchin. *Differential algebra and algebraic groups*. Academic Press, New York, 1973.

[LLL82] A.K. Lenstra, H.W. Lenstra, and L. Lovász. Factoring polynomials with rational coefficients. *Math. Ann.*, 261:515–534, 1982.

[Mas88] D. Masser. Linear relations on algebraic groups. In A. Baker, editor, *New advances in transendence theory*, pages 248–262. Cambridge Univ. Press, 1988.

[Men96] F. Menous. *Les bonnes moyennes uniformisantes et leur application à la resommation réelle*. PhD thesis, Université Paris-Sud, France, 1996.

[Mit96] C. Mitschi. Differential Galois groups of confluent generalized hypergeometric equations: an approach using stokes multipliers. *Pac. J. Math.*, 176(2):365–405, 1996.

[MO95] Scott H. Murray and E. A. O'Brien. Selecting base points for the Schreier-Sims algorithm for matrix groups. *J.S.C.*, 18:577–584, 1995.

[Moe73] R. Moenck. Fast computation of gcds. In *Proc. of the 5th ACM Annual Symposium on Theory of Computing*, pages 142–171, New York, 1973. ACM Press.

[MR91] J. Martinet and J.-P. Ramis. Elementary acceleration and multisummability. *Ann. Inst. Henri Poincaré*, 54(4):331–401, 1991.

[PW02] Victor Y. Pan and Xinmao Wang. Acceleration of euclidean algorithm and extensions. In Teo Mora, editor, *Proc. ISSAC '02*, pages 207–213, Lille, France, July 2002.

[Ram85] J.-P. Ramis. Phénomène de Stokes et filtration Gevrey sur le groupe de Picard-Vessiot. *Notes aux CRAS*, 301(I/5):165–167, 1985.

[Rit50] J.F. Ritt. *Differential algebra*. Amer. Math. Soc., New York, 1950.

[Sch95] L. Schlesinger. *Handbuch der Theorie der linearen Differentialgleichungen*, volume I. Teubner, Leipzig, 1895.

[Sch97] L. Schlesinger. *Handbuch der Theorie der linearen Differentialgleichungen*, volume II. Teubner, Leipzig, 1897.

[Sim70] C.C. Sims. *Computational problems in abstract algebra*, chapter Computational methods in the study of permutation groups, pages 169–183. Pergamon press, Oxford, 1970.

[Sim71] C.C. Sims. Determining the conjugacy classes of permutation groups. In G. Birkhoff and M. Hall Jr., editors, *Computers in algebra and number theory*, volume 4 of *Proc. A.M.S.*, pages 191–195, New York, 1971. A.M.S.

[SU93] M. Singer and F. Ulmer. Galois groups of second and third order linear differential equations. *J.S.C.*, 16:9–36, 1993.

[vdH99] J. van der Hoeven. Fast evaluation of holonomic functions. *TCS*, 210:199–215, 1999.

[vdH01a] J. van der Hoeven. Complex transseries solutions to algebraic differential equations. Technical Report 2001-34, Univ. d'Orsay, 2001.

[vdH01b] J. van der Hoeven. Fast evaluation of holonomic functions near and in singularities. *JSC*, 31:717–743, 2001.

[vdH02] J. van der Hoeven. Relax, but don't be too lazy. *JSC*, 34:479–542, 2002.

[vdH03] J. van der Hoeven. Majorants for formal power series. Technical Report 2003-15, Université Paris-Sud, Orsay, France, 2003.

[vdH04] J. van der Hoeven. Computations with effective real numbers. Technical Report 2004-02, Université Paris-Sud, Orsay, France, 2004. To appear in TCS.

[vdH05a] J. van der Hoeven. Effective complex analysis. *J.S.C.*, 39:433–449, 2005.

[vdH05b] J. van der Hoeven. Efficient accelero-summation of holonomic functions. Technical Report 2005-54, Université Paris-Sud, Orsay, France, 2005. Submitted to JSC.

[vdHea05] J. van der Hoeven et al. Mmxlib: the standard library for Mathemagix, 2002–2005. `http://www.mathemagix.org/mmxweb/web/welcome-mml.en.html`.

[vdP95] M. van der Put. Differential equations modulo *p*. *Compositio Mathematica*, 97:227–251, 1995.

[vdPS03] M. van der Put and M. Singer. *Galois Theory of Linear Differential Equations*, volume 328 of *Grundlehren der mathematischen Wissenschaften*. Springer, 2003.

[vH96] M. van Hoeij. *Factorization of linear differential operators*. PhD thesis, Univ. of Nijmegen, The Netherlands, 1996.

[vH97] M. van Hoeij. Factorization of differential operators with rational functions coefficients. *J.S.C.*, 24:537–561, 1997.

[vHW97] M. van Hoeij and J.A. Weil. An algorithm for computing invariants of differential Galois groups. *J. Pure Appl. Algebra*, 117–118:353–379, 1997.