Uniformly fast evaluation of holonomic functions

JORIS VAN DER HOEVEN

Laboratoire d'informatique, UMR 7161 CNRS
Campus de l'École polytechnique
1, rue Honoré d'Estienne d'Orves
Bâtiment Alan Turing, CS35003
91120 Palaiseau

November 30, 2017

In a series of previous articles, we have given efficient algorithms for the evaluation of holonomic functions over the algebraic numbers and for the computation of their limits at singularities. The focus of these articles was mainly on the efficient evaluation at a fixed point. In the present note, we will show that there exist uniformly efficient algorithms for evaluating holonomic functions. The main technical difficulty is to maintain uniform efficiency near irregular singularities. We will introduce a variant of accelerato-summation for this purpose that we call "expedito-summation".

KEYWORDS: holonomic function, special function, fast evaluation, accelero-summation, expedito-summation

A.M.S. SUBJECT CLASSIFICATION: 33-04, 30-04, 40-04, 33F05, 33E30, 40G10, 30B40

1. INTRODUCTION

Statement of the problem and the main result

Let $\mathbb{K}$ be a subfield of $\mathbb{C}$. A *holonomic function* over $\mathbb{K}$ is a solution f to a linear differential equation $L f = 0$, where $L = \partial^r + L_{r-1} \partial^{r-1} + \cdots + L_0 \in \mathbb{K}(z)[\partial]$ is a monic linear differential operator of order r. Many classical special functions, such as exp, log, sin, cos, erf, hypergeometric functions, Bessel functions, the Airy function, etc. are holonomic. Moreover, the class of holonomic functions is stable under many operations, such as addition, multiplication, differentiation, integration and postcomposition with algebraic functions.

In the sequel, and unless stated otherwise, we will assume that $\mathbb{K}$ is the field of algebraic numbers. The only singularities of a holonomic function f as above can occur at the poles of the rational functions $L_0, \ldots, L_{r-1}$; let Σ denote the finite set of these poles. We will say that f has *initial conditions in* $\mathbb{K}$ if $(f(z), \ldots, f^{(r-1)}(z)) \in \mathbb{K}^r$ for a certain non-singular point $z \in \mathbb{K} \setminus \Sigma$. In this paper, we are interested in the design of efficient algorithms for the numeric evaluation of such a function f, with a particular focus on high precision and uniform efficiency as a function of the argument z.

For a fixed non singular evaluation point, say $z \in \mathbb{K} \setminus \Sigma$, an efficient general purpose algorithm was first given by the Chudnovsky brothers [3]. More precisely, in the case when $\mathbb{K} = \mathbb{Q}[\mathrm{i}]$, they proved that an n-bit approximation of $f(z)$ can be computed in time $O(\mathsf{I}(n) \log^2 n)$. Here $\mathsf{I}(n)$ stands for a complexity bound for integer multiplication and it has recently been proved that one may take $\mathsf{I}(n) = O(n \log n\, 8^{\log^* n})$, where $\log^* n = \min \left\{ k \in \mathbb{N} : \left(\log \circ \overset{k\times}{\cdots} \circ \log\right)(n) \leqslant 1 \right\}$. The Chudnovsky–Chudnovsky algorithm was rediscovered in [7] and generalized to the case when $\mathbb{K}$ is the field of algebraic numbers. An early precursor and further variants can be found in [2, 10, 6].

In order to design uniformly efficient evaluation algorithms, it is crucial to control the efficiency when z approaches one of the singularities in Σ. Actually, one first question concerns the computation of the limit of a holonomic function at a singularity if this limit exists. This was first done in [8] for so called regular singularities (achieving the same complexity bound as for non singular points), and in [9] for irregular singularities (in which case we showed that n-bit approximations of limits can be computed in time $O(\mathsf{I}(n)\log^3 n)$). We refer to [8, 9] for the definitions of the concepts of regular and irregular singularities.

The main aim of this paper is to achieve the same kind of complexity bounds uniformly in z. Such bounds need to be stated with a lot of care. First of all, a holonomic function such as $f(z) = \exp z$ grows exponentially fast at infinity: given the n-bit number $z = 2^n$, one needs $\Theta(2^n)$ bits to merely write down the closest integer approximation $\lfloor \exp(2^n) + 1/2 \rfloor$ of $f(z)$. Using floating point approximations for both z and $f(z)$ does not help, since a similar explosion then occurs for the exponent. But we may hope for a good uniform complexity bound if we use fixed point approximations for z and floating point approximations for $f(z)$.

Another complication is due to the number zero, which should be regarded as a singularity when using floating point representations: it is difficult to compute accurate floating point approximations for $f(z)$ if z is close to a zero of f. Predicting the exact locations of zeros of holonomic functions is a notoriously difficult problem. Even the basic question to decide whether $f(z) = 0$ for $z \in \mathbb{K} \setminus \Sigma$ admits no algorithmic answer for the moment. Nevertheless, the number z is often the approximation of some other complex number with a precision of n binary digits behind the dot. In that case, it is natural to consider the more general evaluation of f on the ball $\mathcal{B}(z, 2^{-n})$ with center z and radius 2^{-n}, and to require that f admits no zeros on this ball.

We are almost in a position to state the main result of this paper. Let $\mathbb{D} = \mathbb{Z}\, 2^{\mathbb{Z}}$ be the set of *dyadic* numbers. Given $x = k\, 2^e \in \mathbb{D}$, we denote by $\operatorname{size}(x) = \lceil \log_2(|k| + 1) \rceil + |e|$ the bitsize of x. Given $z = x + \mathrm{i}\, y \in \mathbb{D}[\mathrm{i}]$, we also denote $\operatorname{size}(z) = \operatorname{size}(x) + \operatorname{size}(y)$. The set $\mathbb{F} = \mathbb{Z}\, 2^{\mathbb{Z}}$ of *floating point* numbers is defined in the same way as $\mathbb{D}$, but the exponent of a floating point numbers $x = k\, 2^e \in \mathbb{F}$ is represented in binary notation, so that the bitsize of x is now fiven by $\operatorname{fsize}(x) = \lceil \log_2(|k| + 1) \rceil + \lceil \log_2(|e| + 1) \rceil$.

Let $z_0 \in \mathbb{K}$ be the point at which we specified the initial conditions of f. We define Ω to be the open subset of $\mathbb{C}$ of all points z such that the straightline segment $[z_0, z]$ from z_0 to z does not intersect Σ. We take f to be the unique solution of $L f = 0$ on Ω that matches the prescribed initial conditions at z_0. Let $\Theta \subseteq \Omega$ denote the set of zeros of f. The main theorem of this paper is the following.

THEOREM 1. *There exists an algorithm that takes $n \in \mathbb{N}$ and $z \in \Omega \cap \mathbb{D}[\mathrm{i}]$ with $\mathcal{B}(z, 2^{-n}) \cap (\partial\Omega \cup \Theta) = \varnothing$ and $\operatorname{size}(z) \leqslant n$ on input and that computes $v \in \mathbb{F}[\mathrm{i}]$ on output with $|f(z) - v| \leqslant 2^{-n}\, |f(z)|$. Moreover, the running time of the algorithm is bounded by $O(\mathsf{I}(n)\log^3 n)$, uniformly in z.*

Proof strategy

As long as z remains in a compact subset K of Ω in Theorem 1, the conclusion essentially follows from the existing complexity bounds in [3, 7]; using a refinement [11] of the complexity analysis from [7], one even obtains the stronger complexity bound $O(\mathsf{I}(n)\log^2 n)$ for the evaluation of f. Using the techniques from [8], these complexity bounds generalize to subsets $K \cap \Omega$ of Ω, where K is a compact set that contains none of the irregular singularities of Σ. If the point at infinity is a regular singularity, then the bound also applies on subsets $\{z \in U : |z| \geqslant M\}$ for sufficiently large M, modulo the change of coordinates $z \to z^{-1}$.

The above discussion shows that the proof of Theorem 1 involves two main difficulties: controlling the complexity near irregular singularities and controlling the complexity of evaluating $f(z)$ near zeros of f. For the first task, we will adapt the technique of accelero-summation from [9]. For the second task, we rely on the idea that $f, \ldots, f^{(r-1)}$ can never simultaneously become "smaller than expected". A precise statement will be presented in Section 4; this statement can be regarded as a quantitative version of the well-known property that $f, \ldots, f^{(r-1)}$ cannot vanish simultaneously unless f vanishes itself.

Let us return to the evaluation of f near an irregular singulary, say $0 \in \Sigma$. At the origin, it is well-known that $L f = 0$ admits a basis of formal solution of the form

$$\tilde{b}_i(z) = \tilde{\varphi}_i(z)\, z^{\lambda_i}\, \mathrm{e}^{P_i(z^{-1/\kappa})}$$

for $L f = 0$, where $\tilde{\varphi}_i(z) \in \mathbb{C}[[z^{1/\kappa}]][\log z]$, $\lambda_i \in \mathbb{C}$, $P_i(z^{-1/\kappa}) \in \mathbb{C}[z^{-1/\kappa}]$, $\kappa \in \mathbb{N}^{\neq}$, and where $\varphi_i(z) \sim (\log z)^{k_i}$ for some $k_i \in \mathbb{N}$. In [9], it is shown that the series $\tilde{\varphi}_i$ are accelero-summable and that we can associate actual functions φ_i to them that are defined on sectors of the form

$$\mathcal{S}_{R,\theta,\alpha} := \{r\, \mathrm{e}^{\vartheta \mathrm{i}} : r \in (0, R], \vartheta \in [\theta - \alpha, \theta + \alpha]\}.$$

Moreover, a finite number of these sectors can be made to cover a punctured neighbourhood of the origin. One crucual step toward the design of an efficient evaluation algorithm for f on such a sector is to deal with the special case when $f = b_i$ for some i, which further reduces to the case when $f = \varphi_i$.

In the remainder of this paper, we will assume that the reader is familiar with [9] and the notations that we used there. For simplicity, we will also restrict to accelerations and Laplace transforms such that we integrate on the positive real axis. Using a change of variables $z \to u\, z$ for a suitable $u \in \mathbb{C}^{\neq}$, this entails no loss of generality. More precisely, we assume that we are in the following situation. The function f is the result

$$f_p \;=\; (\check{\mathcal{L}}_{k_p} \circ \check{\mathcal{A}}_{k_{p-1},k_p} \circ \cdots \circ \check{\mathcal{A}}_{k_1,k_2} \circ \check{\mathcal{B}}_{k_1})(\tilde{f}_1)$$

of an accelero-summation process with critical times $z_1 = \sqrt[k_1]{z}, \ldots, z_p = \sqrt[k_p]{z}$, $k_1 > \cdots > k_p$, and all integrals taken on the positive real axis. The accelero-sum f is defined in some sector

$$\mathcal{S}_{R,\alpha} := \mathcal{S}_{R,0,\alpha}$$

for any $\alpha > 0$ with $\alpha < k_p \pi / 2$.

For any fixed $z \in \mathbb{D}[\mathrm{i}] \cap \mathcal{S}_{R,\alpha}$, the accelero-summation process from [9] provides us with an algorithm to compute n digits of $f(z)$ in time $O(\mathsf{I}(n) \log^3 n)$. In Section 3.1, we will show that this complexity is uniform in z, provided that $n^{k_p} \geqslant \beta_p / |z|$ for some computable constant $\beta_p > 0$. In other words: accelero-summation is a good numerical scheme under the condition that we really need a lot of digits. In [9], we also showed that the technique of "summation until the least term" [12] allows to compute n digits of $f(z)$ in time $O(\mathsf{I}(n) \log^2 n)$, provided that $n^{k_1} \leqslant \beta_1 / |z|$ for some computable constant $\beta_1 > 0$. This complexity bound is also uniform in z.

The above uniform complexity bounds still leave a gap for precisions n between $(\beta_1 / |z|)^{1/k_1}$ and $(\beta_p / |z|)^{1/k_p}$. In order to fill this gap, we introduce the technique of expedito-summation in Section 2. Roughly speaking, we perform the accelero-summation process until some critical time z_q with $1 \leqslant q < p$ and then "expedite" the process by directly taking a truncated Laplace transform with respect to ζ_q. We will show in Section 3.3 that there exist computable constants $\beta_2, \ldots, \beta_{p-1}$ such that expedito-summation until the critical time z_q allows us compute n digits of $f(z)$ in time $O(\mathsf{I}(n) \log^3 n)$, uniformly in z provided that $(\beta_q / |z|)^{1/k_q} \leqslant n \leqslant (\beta_{q+1} / |z|)^{1/k_{q+1}}$.

Notational conventions

This paper should be regarded as a supplement to [9]. For this reason, and as we already stressed before, we will freely use concepts and notations from that paper. In this area it also frequently happens that there exist algorithms to explicitly compute various constants involved in error bounds, but that the precise values of these constants are irrelevant. In [9], we strived to make all error bounds as explicit as possible, but in this paper we will simply denote strictly positive constants of this kind by $\square$. In analysis, the habit to write $O(1)$ for "some bounded function" is somewhat analoguous. For instance, given a real function f and a constant $\sigma \in \mathbb{Q}$, saying that

$$|f(x)| \leqslant \square \, \mathrm{e}^{\square x}$$

for all $x \geqslant \sigma$ means that we can compute an explicit exponential bound for $f(x)$ on the interval $[\sigma, \infty)$.

Acknowledgments. We wish to like Grégoire LECERF for some helpful remarks.

2. EXPEDITO-SUMMATION

Throughout this and the next section, we make the following assumptions:

- $\tilde{f} \in \mathbb{K}[[z^{1/\kappa}]][\log z]$ with $\kappa \in \{1, 2, \ldots\}$ is a formal solution to $L \tilde{f} = 0$.
- $\tilde{f}$ is accelero-summable with critical times $z_1 = \sqrt[k_1]{z}, \ldots, z_p = \sqrt[k_p]{z}$ and $k_1 > \cdots > k_p$.
- The holonomic equations satisfied by the Borel counterparts $\hat{f}_1, \ldots, \hat{f}_p$ at the various critical times admit no singularities on the positive real axis.
- All acceleration integrals and the final Laplace transforms are performed on the positive real axis.

2.1. Introduction to expedito-summation

In [9], we provided a detailed analysis of two summation methods of $\tilde{f}$. The usual accelero-summation process associates the accelero-sum $f = \operatorname{accsum} \tilde{f}$ to $\tilde{f}$ using

$$\operatorname{accsum} \tilde{f} = (\check{\mathcal{L}}_{k_p} \circ \check{\mathcal{A}}_{k_{p-1}, k_p} \circ \cdots \circ \check{\mathcal{A}}_{k_1, k_2} \circ \check{\mathcal{B}}_{k_1})(\tilde{f}).$$

In the appendix, we also considered "summation up to the least term": given $N \in \mathbb{N}$, one may approximate $\operatorname{accsum} \tilde{f}$ by $\operatorname{sum}_N \tilde{f}$, where

$$(\operatorname{sum}_N \tilde{f})(z) = \tilde{f}_0 + \cdots + \tilde{f}_N z^N.$$

Taking $N = (\square |z|)^{-1/k_1}$, we proved that

$$|(\operatorname{sum}_N \tilde{f} - \operatorname{accsum} \tilde{f})(z)| \leqslant \square \, \mathrm{e}^{-(\square |z|)^{-1/k_1}},$$

for all $z \in (0, \square]$.

Summation up to the least term completely shortcuts the whole accelero-summation process. It provides approximations of a precision that correspond to stopping the accelero-summation process at the first singularity for the first critical time. It is natural to consider more general shortcuts, where we perform the usual accelero-summation process up till a given critical time z_q and then "expedite" the remainer of the process by directly performing a truncated Laplace transform on $\zeta_q \in (0, \mathrm{Z}_q]$ for a suitable $\mathrm{Z}_q \in \mathbb{R}^{>}$. More precisely, given $q \in \{1, \ldots, p-1\}$ and $\mathrm{Z}_q \in \mathbb{R}^{>}$, we define

$$\begin{aligned}
(\operatorname{exsum}_{q, \mathrm{Z}_q} \tilde{f})(z_q) &= [(\check{\mathcal{L}}_{k_q, \mathrm{Z}_q} \circ \check{\mathcal{A}}_{k_{q-1}, k_q} \circ \cdots \circ \check{\mathcal{A}}_{k_1, k_2} \circ \check{\mathcal{B}}_{k_1})(\tilde{f}))(z_q) \\
(\check{\mathcal{L}}_{k_q, \mathrm{Z}_q} \check{f}_{k_q})(z_q) &= \int_{\mathcal{H}_{\mathrm{Z}_q}} \check{f}_q(\zeta_q) \, \mathrm{e}^{-\zeta_q / z_q} \, \mathrm{d} \zeta_q.
\end{aligned}$$

Here $\mathcal{H}_{\mathrm{Z}_q}$ denotes the contour from Z_q to $\varepsilon > 0$, turning around 0 and then back from ε to Z_q.

As for summation to the least term, it is natural to chose Z_q such that $\left|\hat{f}_q(\zeta_q)\,\mathrm{e}^{-\zeta_q/z_q}\right|$ is minimal. Since $\hat{f}_q(\zeta_q)$ satisfies a bound of the form

$$\left|\hat{f}_q(\zeta_q)\right| \;\leqslant\; \Box\,\mathrm{e}^{\Box\zeta_q^{\frac{k_q}{k_q-k_{q+1}}}}$$

at infinity, this means that we should take $\Box \leqslant \mathrm{Z}_q \leqslant \mathrm{Z}_{\mathrm{opt}}$, where

$$\mathrm{Z}_{\mathrm{opt}} \;=\; \Box\,|z_q|^{-\frac{k_q-k_{q+1}}{k_{q+1}}} = \Box\,|z_{q+1}|^{-\frac{k_q-k_{q+1}}{k_q}}.$$

Our main aim is to prove the error bound

$$|(\mathrm{exsum}_{q,\mathrm{Z}_q}\tilde{f} - \mathrm{accsum}\,\tilde{f})(z)| \;\leqslant\; \Box\,\mathrm{e}^{-\Box\mathrm{Z}_q/|z_q|} + \Box\,\mathrm{e}^{-\Box/|z_{q+1}|},$$

for $z \in (0, \Box]$. When taking $\mathrm{Z}_q = \mathrm{Z}_{\mathrm{opt}}$, this bound further simplifies to

$$|(\mathrm{exsum}_{q,\mathrm{Z}_q}\tilde{f} - \mathrm{accsum}\,\tilde{f})(z)| \;\leqslant\; \Box\,\mathrm{e}^{-\Box/|z_{q+1}|}.$$

2.2. The expedited approximation

The truncated Laplace transform. Let $\check{g}_q(\zeta_q) = \check{f}_q(\zeta_q)$ for $\operatorname{Re}\zeta_q \leqslant \mathrm{Z}_q$ and $\check{g}_q(\zeta_q) = 0$ for $\operatorname{Re}\zeta_q > \mathrm{Z}_q$, so that

$$g(z) := g_q(z_q) \;:=\; (\check{\mathcal{L}}_{k_q}\check{g}_q)(z_q) = (\check{\mathcal{L}}_{k_q,\mathrm{Z}_q}\check{f}_q)(z_q).$$

Since we know how to compute a bound for $|\check{g}|$ on the contour $\mathcal{H}_{\mathrm{Z}_q}$, we may compute an explicit bound of the form

$$|g_q(z_q)| \;\leqslant\; \Box\,\mathrm{e}^{\Box/|z_q|} \tag{1}$$

for z_q on a small positive sector near zero.

Borel transforms of g at other critical times. For $i = q+1, \ldots, p$, we define

$$\hat{g}_i(\zeta_i) \;:=\; \left(\hat{\mathcal{A}}_{k_q,k_i}\check{g}_q\right)(\zeta_i) = \left(\hat{\mathcal{A}}_{k_q,k_i,\mathrm{Z}_q}\check{f}_q\right)(\zeta_i),$$

where

$$\left(\hat{\mathcal{A}}_{k_q,k_i,\mathrm{Z}_q}\check{f}_q\right)(\zeta_i) \;=\; \int_{\mathcal{H}_{\mathrm{Z}_q}} \check{f}_q(\zeta_q)\,\hat{\mathcal{A}}_{k_q,k_i}(\zeta_q,\zeta_i)\,\mathrm{d}\zeta_q.$$

We may also represent $\hat{g}_i$ as the analytic Borel transform of $g_i(z_i) = g_q(z_q)$ with respect to z_i. Using the bound (1), this allows us to compute a bound

$$|\hat{g}_i(\zeta_i)| \;\leqslant\; \Box\,\mathrm{e}^{\Box\zeta_i} \tag{2}$$

for $\zeta_i \in [\Box, \infty)$.

The difference between f and g. Let $\delta := g - f$. For $i = q, \ldots, p$, we also define

$$\begin{aligned}\hat{\delta}_i &:= \hat{g}_i - \hat{f}_i\\ \delta_i &:= g_i - f_i.\end{aligned}$$

For $i = q$, and setting $c_i = \mathrm{Z}_q$, we thus have

$$\hat{\delta}_i(\zeta_i) \;=\; 0$$

for $\zeta_i \in (0, c_i]$. One major topic of this section will be to compute bounds at the origin for $\left|\hat{\delta}_i(\zeta_i)\right|$ and $i > q$.

Behaviour of the $\hat{\delta}_i$ at infinity. Let $q \leqslant i < p$. Combining the bound (2) with the superexponential bound for $\hat{f}_i$, as provided by the accelero-summation process, we may compute a bound

$$\left|\hat{\delta}_i(\zeta_i)\right| \leqslant \square \, \mathrm{e}^{\square \zeta_i^{\frac{k_i}{k_i - k_{i+1}}}} \tag{3}$$

for $\zeta_i \in [\square, \infty)$. Notice that we may also compute a bound

$$\left|\hat{\mathcal{A}}_{k_i,k_{i+1}}(\zeta_i)\right| \leqslant \square \, \mathrm{e}^{-\square\left(\zeta_i^{\frac{k_i}{k_i-k_{i+1}}} / \zeta_{i+1}^{\frac{k_{i+1}}{k_i-k_{i+1}}}\right)} \tag{4}$$

for $\zeta_i \in [\square, \infty)$ and $\zeta_{i+1} \in (0, \square]$.

Majorants for specific accelerates and Laplace transforms. The following bounds will be useful for proving precise error estimates for the $\hat{\delta}_i$ and δ_p. The proofs are a routine application of the saddle point technique.

LEMMA 2. *Let z_q and z_i be critical times with $q < i$. Then*

$$\left|\int_0^\infty \mathrm{e}^{-\square \zeta_i^{\frac{-k_i}{k_q-k_i}}} \mathrm{e}^{-\square \zeta_i/|z_i|} \, \mathrm{d}\zeta_i\right| \leqslant \square \, \mathrm{e}^{-\square/|z_q|}, \tag{5}$$

for all $|z_q| \in (0, \square]$. If $q+1 < i$, then

$$\left|\int_0^\infty \mathrm{e}^{-\square \zeta_{i-1}^{\frac{-k_{i-1}}{k_q-k_{i-1}}}} \hat{\mathcal{A}}_{k_{i-1},k_i}(\zeta_{i-1}, \zeta_i) \, \mathrm{d}\zeta_{i-1}\right| \leqslant \square \, \mathrm{e}^{-\square \zeta_i^{\frac{-k_i}{k_q-k_i}}}, \tag{6}$$

for all $\zeta_i \in (0, \square]$.

2.3. The first acceleration

LEMMA 3. *Let $i = q+1$. We can compute $c_i > 0$ such that, for $\zeta_i \in (0, c_i]$, we have*

$$\left|\hat{\delta}_i(\zeta_i)\right| \leqslant \square \, \mathrm{e}^{-\square \mathrm{Z}_q^{\frac{k_q}{k_q-k_i}} \zeta_i^{\frac{-k_i}{k_q-k_i}}}.$$

Proof. We have

$$\begin{aligned} \hat{\delta}_i(\zeta_i) &= (\hat{\mathcal{A}}_{k_q,k_i}(\hat{\delta}_q))(\zeta_i) \\ &= \int_{\mathrm{Z}_q}^\infty \hat{\delta}_q(\zeta_q) \, \hat{\mathcal{A}}_{k_q,k_i}(\zeta_q, \zeta_i) \, \mathrm{d}\zeta_q. \end{aligned}$$

Using (3) and (4), it follows that

$$\begin{aligned} \left|\hat{\delta}_i(\zeta_i)\right| &\leqslant \left|\square \int_{\mathrm{Z}_q}^\infty \mathrm{e}^{\square \zeta_q^{\frac{k_q}{k_q-k_i}} - \square\left(\zeta_q^{\frac{k_q}{k_q-k_i}} / \zeta_i^{\frac{k_i}{k_q-k_i}}\right)} \mathrm{d}\zeta_q\right| \\ &\leqslant \square \, \mathrm{e}^{-\square \mathrm{Z}_q^{\frac{k_q}{k_q-k_i}} \zeta_i^{\frac{-k_i}{k_q-k_i}}}, \end{aligned}$$

on an interval $\zeta_i \in (0, c_i]$ for some computable $c_i > 0$. $\square$

2.4. Subsequent accelerations

LEMMA 4. *For each $i > q+1$, we can compute a constant $c_i > 0$, together with a bound*

$$\left|\hat{\delta}_i(\zeta_i)\right| \leqslant \square\, \mathrm{e}^{-\square \mathrm{Z}_q^{\frac{k_q}{k_q-k_i}} \zeta_i^{\frac{-k_i}{k_q-k_i}}} + \square\, \mathrm{e}^{-\square \zeta_i^{\frac{-k_i}{k_{q+1}-k_i}}}$$

for $\zeta_i \in (0, c_i]$.

Proof. We will prove the lemma by induction over i. We have

$$\hat{\delta}_i(\zeta_i) = (\hat{\mathcal{A}}_{k_{i-1},k_i}(\hat{\delta}_{i-1}))(\zeta_i) = I_1(\zeta_i) + I_2(\zeta_i),$$

where

$$\begin{aligned} I_1(\zeta_i) &:= \int_0^{c_{i-1}} \hat{\delta}_{i-1}(\zeta_{i-1})\, \hat{\mathcal{A}}_{k_{i-1},k_i}(\zeta_{i-1}, \zeta_i)\, \mathrm{d}\,\zeta_{i-1} \\ I_2(\zeta_i) &:= \int_{c_{i-1}}^{\infty} \hat{\delta}_{i-1}(\zeta_{i-1})\, \hat{\mathcal{A}}_{k_{i-1},k_i}(\zeta_{i-1}, \zeta_i)\, \mathrm{d}\,\zeta_{i-1} \end{aligned}$$

If $i = q+2$, then Lemma 3 yields a bound

$$\left|\hat{\delta}_{i-1}(\zeta_{i-1})\right| \leqslant \square\, \mathrm{e}^{-\square \mathrm{Z}_q^{\frac{k_q}{k_q-k_{i-1}}} \zeta_{i-1}^{\frac{-k_{i-1}}{k_q-k_{i-1}}}},$$

for $\zeta_{i-1} \in (0, c_{i-1}]$. For $i > q+2$, the induction hypothesis yields the bound

$$\left|\hat{\delta}_{i-1}(\zeta_{i-1})\right| \leqslant \square\, \mathrm{e}^{-\square \mathrm{Z}_q^{\frac{k_q}{k_q-k_{i-1}}} \zeta_{i-1}^{\frac{-k_{i-1}}{k_q-k_{i-1}}}} + \square\, \mathrm{e}^{-\square \zeta_{i-1}^{\frac{-k_{i-1}}{k_{q+1}-k_{i-1}}}},$$

for $\zeta_{i-1} \in (0, c_{i-1}]$. Now, in view of (6), we may compute a suffiently small $c_i > 0$ such that

$$\left|\int_0^{c_{i-1}} \mathrm{e}^{-\square \mathrm{Z}_q^{\frac{k_q}{k_q-k_{i-1}}} \zeta_{i-1}^{\frac{-k_{i-1}}{k_q-k_{i-1}}}}\, \hat{\mathcal{A}}_{k_{i-1},k_i}(\zeta_{i-1}, \zeta_i)\, \mathrm{d}\,\zeta_{i-1}\right| \leqslant \square\, \mathrm{e}^{-\square \mathrm{Z}_q^{\frac{k_q}{k_q-k_i}} \zeta_i^{\frac{-k_i}{k_q-k_i}}},$$

for all $\zeta_i \in (0, c_i]$. If $i > q+2$, then a similar computation yields

$$\left|\int_0^{c_{i-1}} \mathrm{e}^{-\square \zeta_{i-1}^{\frac{-k_{i-1}}{k_{q+1}-k_{i-1}}}}\, \hat{\mathcal{A}}_{k_{i-1},k_i}(\zeta_{i-1}, \zeta_i)\, \mathrm{d}\,\zeta_{i-1}\right| \leqslant \square\, \mathrm{e}^{-\square \zeta_i^{\frac{-k_i}{k_{q+1}-k_i}}},$$

for all $\zeta_i \in (0, c_i]$, modulo a decrease of c_i if necessary. Putting these bounds together, we obtain

$$|I_1(\zeta_i)| \leqslant \square\, \mathrm{e}^{-\square \mathrm{Z}_q^{\frac{k_q}{k_q-k_i}} \zeta_i^{\frac{-k_i}{k_q-k_i}}} + \square\, \mathrm{e}^{-\square \zeta_i^{\frac{-k_i}{k_{q+1}-k_i}}}$$

for $\zeta_i \in (0, c_i]$. Using (3) and (4), we may also compute a bound

$$\begin{aligned} |I_2(\zeta_i)| &\leqslant \left|\square \int_{c_{i-1}}^{\infty} \mathrm{e}^{\square \zeta_{i-1}^{\frac{k_{i-1}}{k_{i-1}-k_i}} - \square\left(\zeta_{i-1}^{\frac{k_{i-1}}{k_{i-1}-k_i}} / \zeta_i^{\frac{k_i}{k_{i-1}-k_i}}\right)}\, \mathrm{d}\,\zeta_{i-1}\right| \\ &\leqslant \square\, \mathrm{e}^{-\square \zeta_i^{\frac{-k_i}{k_{i-1}-k_i}}} \\ &\leqslant \square\, \mathrm{e}^{-\square \zeta_i^{\frac{-k_i}{k_{q+1}-k_i}}} \end{aligned}$$

for $\zeta_i \in (0, c_i]$, modulo a further decrease of c_i if necessary, and where we used the fact that $k_{q+1} \geqslant k_{i-1}$. Combining the bounds for I_1 and I_2, the result follows. □

2.5. The final Laplace transform

LEMMA 5. *For any aperture* $\alpha \in (0, \pi/2) \cap \mathbb{Q}$, *we can compute a* $\sigma > 0$ *and a bound*

$$|\delta_p(z_p)| \;\leqslant\; \square\, \mathrm{e}^{-\square \mathrm{Z}_q / |z_q|} + \square\, \mathrm{e}^{-\square / |z_{q+1}|}$$

for all $z_p \in \mathcal{S}_{\sigma,\alpha}$.

Proof. We have

$$\delta_p(z_p) \;=\; (\hat{\mathcal{L}}_{k_p}(\hat{\delta}_p))(z_p) = I_1(z_p) + I_2(z_p),$$

where

$$\begin{aligned} I_1(z_p) &:= \int_0^{c_p} \hat{\delta}_p(\zeta_p)\, \mathrm{e}^{-\zeta_p / z_p}\, \mathrm{d}\, \zeta_p \\ I_2(z_p) &:= \int_{c_p}^{\infty} \hat{\delta}_p(\zeta_p)\, \mathrm{e}^{-\zeta_p / z_p}\, \mathrm{d}\, \zeta_p. \end{aligned}$$

Using Lemma 4 and (5), we can compute $\sigma > 0$ and a bound

$$\begin{aligned} |I_1(z_p)| &\leqslant \int_0^{c_p} \left(\square\, \mathrm{e}^{-\square \mathrm{Z}_q^{\frac{k_q}{k_q - k_p}} \zeta_p^{\frac{-k_p}{k_q - k_p}}} + \square\, \mathrm{e}^{-\square \zeta_p^{\frac{-k_p}{k_{q+1} - k_p}}} \right) |\mathrm{e}^{-\zeta_p / z_p}|\, \mathrm{d}\, \zeta_p \\ &\leqslant \int_0^{c_p} \left(\square\, \mathrm{e}^{-\square \mathrm{Z}_q^{\frac{k_q}{k_q - k_p}} \zeta_p^{\frac{-k_p}{k_q - k_p}}} + \square\, \mathrm{e}^{-\square \zeta_p^{\frac{-k_p}{k_{q+1} - k_p}}} \right) \mathrm{e}^{-\square \zeta_p / |z_p|}\, \mathrm{d}\, \zeta_p \\ &\leqslant \square\, \mathrm{e}^{-\square \mathrm{Z}_q / |z_q|} + \square\, \mathrm{e}^{-\square / |z_{q+1}|} \end{aligned}$$

for $z_p \in \mathcal{S}_{\sigma,\alpha}$. Using (2) and the exponential bound for $\hat{f}_p$ as provided by the accelero-summation process, we may also compute a bound

$$|\hat{\delta}(\zeta_p)| \;\leqslant\; \square\, \mathrm{e}^{\square \zeta_p}$$

for $\zeta_p \in [c_p, \infty)$. Modulo a further increase of σ if necessary, this allows us to compute a bound

$$\begin{aligned} |I_2(z_p)| &\leqslant \int_{c_p}^{\infty} \square\, \mathrm{e}^{\square \zeta_p}\, |\mathrm{e}^{-\zeta_p / z_p}|\, \mathrm{d}\, \zeta_p \\ &\leqslant \int_{c_p}^{\infty} \square\, \mathrm{e}^{\square \zeta_p}\, \mathrm{e}^{-\square \zeta_p / |z_p|}\, \mathrm{d}\, \zeta_p \\ &\leqslant \square\, \mathrm{e}^{-\square / |z_p|} \\ &\leqslant \square\, \mathrm{e}^{-\square / |z_q|}, \end{aligned}$$

for $z_p \in \mathcal{S}_{\sigma,\alpha}$. Adding up the bounds for I_1 and I_2, the results follows. □

COROLLARY 6. *For any aperture* $\alpha \in (0, \pi/2) \cap \mathbb{Q}$, *and assuming that*

$$\square \leqslant \mathrm{Z}_q \leqslant \square\, |z_q|^{-\frac{k_q - k_{q+1}}{k_{q+1}}},$$

we can compute a $\sigma > 0$ *and a bound*

$$|\delta_p(z_p)| \;\leqslant\; \square\, \mathrm{e}^{-\square \mathrm{Z}_q / |z_q|}$$

for all $z_p \in \mathcal{S}_{\sigma,\alpha}$.

Proof. This directly follows from the fact that $z_q^{-\frac{k_q - k_{q+1}}{k_{q+1}}} / z_q = z_{q+1}$. □

3. UNIFORM COMPLEXITY ON LOCAL SECTORS

3.1. Uniform complexity of accelero-summation

PROPOSITION 7. *Let $\alpha \in (0, \pi/2) \cap \mathbb{D}$ be a fixed aperture and let $\beta_p \in \mathbb{D}^>$. Then we can compute a constant $\sigma > 0$ with the following property: given $z \in \mathbb{D}[\mathrm{i}]$ and $n \geqslant \mathrm{size}(z) + \square$ with*

$$\beta_p / n^{k_p} \leqslant |z| \leqslant \sigma$$

and $|\arg z| \leqslant k_p\, \alpha$, we can compute an approximation $v \in \mathbb{D}[\mathrm{i}]$ with $|v - f(z)| \leqslant 2^{-n}$ in time $O(\mathsf{I}(n) \log^3 n)$, where the complexity bound holds uniformly in z under the above conditions.

Proof. Recall that we may compute an exponential bound

$$|\hat{f}_p(\zeta_p)| \;\leqslant\; \square\, \mathrm{e}^{\square \zeta_p}$$

for $\hat{f}_p$ at infinity. For $\mathrm{Z}_p = \square\, n\, |z_p|$ and $n \geqslant \square$, this yields a bound

$$\int_{\mathrm{Z}_p}^{\infty} |\hat{f}_p(\zeta_p)|\, |\mathrm{e}^{-z_p/\zeta_p}|\, \mathrm{d}\zeta_p \;\leqslant\; 2^{-n-1}.$$

We now wish to compute v by approximating the truncated Laplace integral

$$u \;:=\; \int_{\mathcal{H}_{\mathrm{Z}_p}} \check{f}_p(\zeta_p)\, \mathrm{e}^{-z_p/\zeta_p}\, \mathrm{d}\zeta_p \tag{7}$$

with precision 2^{-n-1}, i.e. $|v - u| \leqslant 2^{-n-1}$ and $|v - f(z)| \leqslant 2^{-n}$.

Let us first consider the case when the bitsize of z_p is bounded by $\square \log n$. Under the assumption that $|z| \geqslant \beta_p / n^{k_p}$, we observe that $\mathrm{Z}_p \geqslant \square$. This implies that we can chose the contour $\mathcal{H}_{\mathrm{Z}_p}$ to use a circle of fixed radius around the origin (which does not depend on z_p). We next evaluate (7) using the algorithm from [9, Section 6]. Our hypothesis that $\mathrm{size}(z_p) = O(\log n)$ implies that the primitive of $\check{f}_p(\zeta_p)\, \mathrm{e}^{-z_p/\zeta_p}$ satisfies a holonomic equation of size $O(\log n)$, uniformly in z_p. Consequently, it can be checked that the complexity bound from [9] holds uniformly in z_p. This means that the required 2^{-n-1}-approximation v of u can be computed in time $O(\mathsf{I}(n) \log^3 n)$, uniformly in z_p.

For general z, we approximate $f(z)$ in two steps. Let $\kappa = k_1$ be the growth rate of the linear differential equation satisfied by f at the origin. In [9, Theorem 5.2], we showed that in the sector $\mathcal{S} = \mathcal{S}_{\square, k_p \alpha}$, we have the following bound for the transition matrix on a straightline path $z \rightarrow z'$ in $\mathcal{S}$:

$$\|\Delta_{z \rightarrow z'}\| \;\leqslant\; \square\, \mathrm{e}^{\square |(z')^{-\kappa} - z^{-\kappa}|}.$$

For $z' - z \leqslant \square\, z^{\kappa+1}$, it follows that

$$\|\Delta_{z \rightarrow z'}\| \;\leqslant\; \square. \tag{8}$$

Now let $z' \in \mathbb{D}[\mathrm{i}]$ be an approximation of z with $z' - z \leqslant \square\, z^{\kappa+1}$ and $\mathrm{size}(z') \leqslant \square\, |\log|z||$. By what precedes, we may compute $2^{-n-\square}$-approximations of $f(z'), \ldots, f^{(r-1)}(z')$ in time $O(\mathsf{I}(n) \log^3 n)$, uniformly in z. Using the usual "bitburst" algorithm from [3, 7, 11], together with (8), it follows that we may compute a 2^{-n}-approximation of $f(z)$ using an additional time of $O(\mathsf{I}(n) \log^2 n)$, uniformly in z. Adding up these complexity bounds, the result follows. □

3.2. Uniform complexity of summation until the least term

PROPOSITION 8. *Let* $\alpha \in (0, \pi/2) \cap \mathbb{D}$. *Then we may compute a constants* $\beta_1 \in \mathbb{D}^>$ *such that*

$$|(\mathrm{sum}_N \tilde{f})(z) - f(z)| \ \leqslant \ 2^{-n-1}$$

for all $z \in \mathbb{C}$, $n \geqslant \square$ *and* $N = \square n$ *with*

$$|z| \leqslant \beta_1 / n^{k_1}$$

and $|\arg z| \leqslant k_p\, \alpha$. *Moreover, if* $z \in \mathbb{D}[\mathrm{i}]$ *and* $n \geqslant \mathrm{size}(z)$, *then we can compute an approximation* $v \in \mathbb{D}[\mathrm{i}]$ *with* $|v - (\mathrm{sum}_N \tilde{f})(z)| \leqslant 2^{-n}$ *in time* $O(\mathsf{I}(n) \log^2 n)$, *where the complexity bound holds uniformly in* z *under the above conditions.*

Proof. Direct consequence of [9, Theorem A.1]. □

3.3. Uniform complexity of expedito-summation

PROPOSITION 9. *Let* $\alpha \in (0, \pi/2) \cap \mathbb{D}$ *be a fixed aperture and let* $\beta_q \in \mathbb{D}^>$, *where* $q < p$. *Then we may compute a constant* $\beta_{q+1} \in \mathbb{D}^>$ *such that*

$$|(\mathrm{exsum}_{q, \mathrm{Z}_q} \tilde{f})(z) - f(z)| \ \leqslant \ 2^{-n-1}$$

for all $z \in \mathbb{C}$ *and* $n \geqslant \square$ *with*

$$\beta_q / n^{k_q} \leqslant |z| \leqslant \beta_{q+1} / n^{k_{q+1}}$$

and $|\arg z| \leqslant k_p\, \alpha$, *where*

$$\mathrm{Z}_q \ = \ \square n\, |z_q|.$$

Moreover, if $z \in \mathbb{D}[\mathrm{i}]$ *and* $n \geqslant \mathrm{size}(z)$, *then we can compute an approximation* $v \in \mathbb{D}[\mathrm{i}]$ *with* $|v - (\mathrm{exsum}_{q, \mathrm{Z}_q} \tilde{f})(z)| \leqslant 2^{-n-1}$ *in time* $O(\mathsf{I}(n) \log^3 n)$, *where the complexity bound holds uniformly in* z *under the above conditions.*

Proof. Our hypothesis on $|z|$ implies that

$$\square \leqslant \mathrm{Z}_q \leqslant \square\, |z_q|^{-\frac{k_q - k_{q+1}}{k_{q+1}}}.$$

By Corollary 6, it follows that for all z_p with $|z| = |z_p^{k_p}| \leqslant \square$ and $|\arg z| = k_p\, |\arg z_p| \leqslant k_p\, \alpha$, we have

$$|(\mathrm{exsum}_{q, \mathrm{Z}_q} \tilde{f})(z) - f(z)| \ = \ |\delta_p(z_p)| \ = \square\, \mathrm{e}^{-\square \mathrm{Z}_q / |z_q|} \ \leqslant \ 2^{-n-1}.$$

For any suitable point ζ_q^{init} close to the origin and $i \in \mathbb{N}$, we have shown in [9] how to compute n decimal digits of $\check{f}_q^{(i)}(\zeta_q^{\mathrm{init}})$ in time $O(\mathsf{I}(n) \log^3 n)$. This provides us with the required initial conditions for the analytic continuation of the integrant of the truncated Laplace integral

$$u \ := \ \int_{\mathcal{H}_{\mathrm{Z}_q}} \check{f}_q(\zeta_q)\, \mathrm{e}^{-z_q / \zeta_q}\, \mathrm{d}\, \zeta_q.$$

In a similar way as in the proof of Proposition 7, we may therefore approximate u to precision 2^{-n-1} in time $O(\mathsf{I}(n) \log^3 n)$, where the complexity bound is uniform in z under our conditions. □

3.4. The combined local strategy

Putting Propositions 7, 8 and 9 together, we obtain:

THEOREM 10. *Let $\alpha \in (0, \pi/2) \cap \mathbb{D}$ be a fixed aperture. Then we may compute a constant $\sigma \in \mathbb{D}^>$ with the following property. Given $z \in \mathbb{D}[\mathrm{i}]^{\neq}$ and $n \in \mathbb{N}$ on input with $|\arg z| \leqslant k_p \alpha$ and $|z| \leqslant \sigma$, we may compute a 2^{-n}-approximation of $f(z)$ in time $O(\mathsf{I}(n) \log^3 n)$, where the complexity bound holds uniformly in z.*

Proof. Let $\sigma, \beta_1, \ldots, \beta_p$ be as in Propositions 8, 9 and 7. For any $z \in \mathbb{C}^{\neq}$ with $|\arg z| \leqslant k_p \alpha$ and $|z| \leqslant \sigma$, at least one of the following three statements holds:

1. We have $\beta_p / n^{k_p} \leqslant |z| \leqslant \sigma$.
2. We have $\beta_q / n^{k_q} \leqslant |z| \leqslant \beta_{q+1} / n^{k_{q+1}}$ for some $q \in \{1, \ldots, p-1\}$.
3. We have $|z| \leqslant \beta_1 / n^{k_1}$.

In these cases we respectively apply Proposition 7, 9 or 8 in order to obtain the desired result. □

4. GLOBALLY EFFICIENT EVALUATION

4.1. Local analysis of cancellations

Assume that L is singular at the origin. Then for some $\kappa \in \mathbb{N}$, there exists a basis of formal solutions of the form

$$\tilde{b}_i(z) = \tilde{\varphi}_i(z)\, z^{\lambda_i}\, \mathrm{e}^{P_i(z^{-1/\kappa})} \tag{9}$$

for $L f = 0$, where $\tilde{\varphi}_i(z) \in \mathbb{C}[[z^{1/\kappa}]][\log z]$, $\lambda_i \in \mathbb{C}$, $P_i(z^{-1/\kappa}) \in \mathbb{C}[z^{-1/\kappa}]$, and where $\varphi_i(z) \sim (\log z)^{k_i}$ for some $k_i \in \mathbb{N}$. Moreover, each $\tilde{\varphi}_i$ belongs to the subset $\mathbb{A}$ of $\mathbb{C}[[z^{1/\kappa}]][\log z]$ accelero-summable series.

For each fixed accelero-summation scheme, there exist ρ, θ and α such that the $\tilde{\varphi}_i(z)$ and $\tilde{b}_i(z)$ give rise to analytic functions $\varphi_{\mathcal{S},i}(z)$ and $b_{\mathcal{S},i}(z)$ on the sector $\mathcal{S} = \mathcal{S}_{\rho,\theta,\alpha}$. A sector $\mathcal{S}$ for which this happens is said to be *admissible*. Moreover, there exist a finite number of admissible sectors $\mathcal{S}_{\rho_1,\theta_1,\alpha_1}, \ldots, \mathcal{S}_{\rho_\ell,\theta_\ell,\alpha_\ell}$ with $\rho_i, \mathrm{e}^{\mathrm{i}\theta_i}, \mathrm{e}^{\mathrm{i}\alpha_i} \in \mathbb{K}$ whose interiors cover a small neighbourhood of $\mathbb{C}^{\neq}$. We will call this an *admissible cover*.

Let $\mathcal{S} = \mathcal{S}_{\rho,\theta,\alpha}$ be one of the sectors in an admissible cover and let φ_i and b_i denote the accelero-sums of $\tilde{\varphi}_i$ and $\tilde{b}_i$ on this sector. For each $i \in \{1, \ldots, r\}$, let $E_i(z) = z^{\sigma_i}\, \mathrm{e}^{P_i(z^{-1/\kappa})}$. Let $\mathcal{S}_{\mathrm{Id}}$ denote the subset of all $z \in \mathcal{S}$ such that

$$|E_1(z)| \geqslant |E_2(z)| \geqslant \cdots \geqslant |E_r(z)|.$$

More generally, given a permutation π of $\{1, \ldots, r\}$, let $\mathcal{S}_\pi$ denote the subset of all $z \in \mathcal{S}$ with

$$|E_{\pi(1)}(z)| \geqslant |E_{\pi(2)}(z)| \geqslant \cdots \geqslant |E_{\pi(r)}(z)|.$$

Clearly, $\mathcal{S} = \bigcup_\pi \mathcal{S}_\pi$.

Let $f = \lambda_1 b_1 + \cdots + \lambda_r b_r$ be a non zero solution to $L f = 0$ on $\mathcal{S}$ and let F be the column vectors with entries f, f', ..., $f^{(r-1)}$. Although f can vanish on $\mathcal{S}$ due to cancellations among the terms $\lambda_i b_i$ and $\lambda_j b_j$, the vector F cannot vanish unless $f = 0$. We will now prove a stronger version of this observation by showing that the sup-norm $\|F\|$ of F cannot become much smaller than $|E_1(z)|$.

THEOREM 11. *There exist constants $C > 0$ and ν such that*

$$\|F(z)\| \ \geqslant\ C\, |E_1(z)\, z^\nu|,$$

for all $z \in \mathcal{S}_{\mathrm{Id}}$.

Proof. Without loss of generality, we may assume that $|z| \leqslant 1$. For each $k \in \{1, ..., r\}$, let W_r be the Wronskian matrix

$$W_k(z) = \begin{pmatrix} b_1(z) & \cdots & b_k(z) \\ \vdots & & \vdots \\ b_1^{(k-1)}(z) & \cdots & b_k^{(k-1)}(z) \end{pmatrix}$$

We may decompose

$$W_k(z) = U_k(z)\,\Delta_k(z),$$

where

$$\Delta_k(z) = \begin{pmatrix} E_1(z) & & \\ & \ddots & \\ & & E_k(z) \end{pmatrix},$$

and where the entries of U_k are in $\mathbb{A}\, z^{-\pi_k}$ for some $\pi_k \in \mathbb{N}$ that only depends on the degrees of $P_1, \ldots, P_k$. It follows that

$$W_k^{-1}(z) = \Delta_k^{-1}(z)\,\frac{\operatorname{adj}(U_k(z))}{\det(U_k(z))},$$

where $\det(U_k(z)) \neq 0$ by the linear independence of $b_1, \ldots, b_k$. Now $\det(U_k(z))$ and the entries of $\operatorname{adj}(U_k(z))$ are all elements of $\mathbb{A}\, z^{-k\pi_k}$. It follows that there exists a constant $\nu_k \in \mathbb{R}$ such that $\operatorname{adj}(U_k(z)) / \det(U_k(z)) = O(z^{-\nu_k})$ for all $z \in \mathcal{S}_{\mathrm{Id}}$.

Now consider our fixed linear combination $f(z) = \lambda_1\, b_1(z) + \cdots + \lambda_r\, b_r(z)$ and let

$$\Lambda_k = \begin{pmatrix} \lambda_1 \\ \vdots \\ \lambda_k \end{pmatrix}, \qquad G_k(z) = \begin{pmatrix} g_k(z) \\ \vdots \\ g_k^{(k-1)}(z) \end{pmatrix},$$

where $g_k(z) = \lambda_1\, b_1(z) + \cdots + \lambda_k\, b_k(z)$, so that $F = G_r$ and $G_k = W_k\, \Lambda_k$. Also let $E(z)$ be the column vector with entries $E_1(z)$, ..., $E_r(z)$. For the sup-norm on vectors, the above discussion shows that

$$\|\Lambda_k\| = O(E_k^{-1}(z)\, z^{-\nu_k}\, \|G_k(z)\|).$$

For some fixed constant $C_k > 0$, this means that

$$\|G_k(z)\| \geqslant C_k\, |E_k(z)\, z^{\nu_k}|. \tag{10}$$

There also exist constants $M > 0$ and μ such that for all $k \in \{1, \ldots, r\}$ and $i < r$,

$$r\, |(\lambda_k\, \varphi_k(z)\, E_k(z))^{(i)}\, E_k(z)^{-1}| \leqslant M\, |z^{\mu}|. \tag{11}$$

Now we may partition $\mathcal{S}_{\mathrm{Id}}$ into r subsets $\mathcal{S}_{\mathrm{Id},1}, \ldots, \mathcal{S}_{\mathrm{Id},r}$ as follows. By induction over k, we define $\mathcal{S}_{\mathrm{Id},k}$ to be the subset of all $z \in \mathcal{S}_{\mathrm{Id}} \setminus (\mathcal{S}_{\mathrm{Id},1} \cup \cdots \cup \mathcal{S}_{\mathrm{Id},k-1})$ such that

$$2\, M\, |E_{k+1}(z)\, z^{\mu}| < C_k\, |E_k(z)\, z^{\nu_k}|,$$

where we understand that $E_{r+1}(z) = 0$. If $z \in \mathcal{S}_{\mathrm{Id},k}$, then it follows that

$$\begin{aligned} |E_2(z)| &\geqslant (C_1 / 2\, M)\, |E_1(z)\, z^{\nu_1 - \mu}| \\ |E_3(z)| &\geqslant (C_1\, C_2 / 4\, M^2)\, |E_1(z)\, z^{\nu_1 + \nu_2 - 2\mu}| \\ &\vdots \\ |E_k(z)| &\geqslant (C_1 \cdots C_{k-1} / 2^{k-1}\, M^{k-1})\, |E_1(z)\, z^{\nu_1 + \cdots + \nu_{k-1} - (k-1)\mu}| \end{aligned}$$

Still for $z \in \mathcal{S}_{\mathrm{Id},k}$, the relation (10) also implies the existence of an $i < k$ such that

$$\left| g_k^{(i)}(z) \right| \;\geqslant\; C_k \, |E_k(z) \, z^{\nu_k}|.$$

Using (11), it follows that

$$\begin{aligned} & 2 \, |(f - g_k)^{(i)}(z)| \\ \leqslant\; & 2 \, M \, |E_{k+1}(z) \, z^{\mu}| \\ <\; & C_k \, |E_k(z) \, z^{\nu_k}| \\ \leqslant\; & \left| g_k^{(i)}(z) \right|, \end{aligned}$$

whence

$$\begin{aligned} |f^{(i)}(z)| \;&\geqslant\; \tfrac{1}{2} \left| g_k^{(i)}(z) \right| \\ &\geqslant\; \tfrac{1}{2} \, C_k \, |E_k(z) \, z^{\nu_k}| \\ &\geqslant\; (C_1 \cdots C_k / 2^k \, M^{k-1}) \, |E_1(z) \, z^{\nu_1 + \cdots + \nu_k - (k-1) \mu}|. \end{aligned}$$

We conclude that $\|F(z)\| \geqslant C \, |E_1(z) \, z^{\nu}|$ for $C = \min \{C_1 \cdots C_k / 2^k \, M^{k-1} : 1 \leqslant k \leqslant r\}$ and $\nu = \max \{\nu_1 + \cdots + \nu_k - (k-1) \, \mu : 1 \leqslant k \leqslant r\}$, using our assumption that $|z| \leqslant 1$. □

Remark 12. It is plausible that a bound for ν can be stated in terms of κ and the degrees of $P_1, \ldots, P_r$. We have not pursued this line of thought any further since any constant ν will do for our purposes.

4.2. Existence of zeros on disks

Consider the power series expansion $f(z+t) = f_0 + f_1 \, t + f_2 \, t^2 + \cdots$ of f at z. For each $k \in \mathbb{N}$, let Φ_k be the vector with entries $f_k, \ldots, f_{k+r-1}$. Theorem 11 provides us with a uniform lower bound for $\|\Phi_0\|$ in terms of E_1. We also have the following upper bound for the remaining coefficients.

LEMMA 13. *There exist constants $\varrho > 0$, $A > 0$ and $\tau \in \mathbb{Z}$ such that*

$$\|\Phi_k\| \;\leqslant\; \|\Phi_0\| \, |A \, z^{\tau}|^k,$$

for all $k \in \mathbb{N}$ and $z \in \mathcal{S}$ with $|z| \leqslant \varrho$.

Proof. Since f is holonomic, there exists a matrix M_k with coefficients in $\mathbb{K}(z)[k^{-1}]$ such that

$$\Phi_{k+1} \;=\; M_k \, \Phi_k.$$

Consequently, there exists a uniform majorant equation for Φ_{k+1} of the form

$$\bar{\Phi}_{k+1} \;=\; \frac{A}{r} \, J \, \bar{\Phi}_k \, |z^{\tau}|,$$

for suitable constants $A > 0$ and $\tau \in \mathbb{Z}$, and where J denotes the $r \times r$ matrix whose coefficients are all one. Taking $\bar{\Phi}_0$ to be the vector with entries $\|\Phi_0\|, \ldots, \|\Phi_0\|$, it follows that $\bar{\Phi}_k$ is the vector with entries $\|\Phi_0\| \, |A \, z^{\tau}|^k, \ldots, \|\Phi_0\| \, |A \, z^{\tau}|^k$. By construction $\|\Phi_k\| \leqslant \|\bar{\Phi}_k\|$. □

LEMMA 14. *Let $g = g_0 + g_1 \, t + \cdots$ be an analytic function on the unit disk $\mathcal{B}(0,1)$ such that $|g_0| \leqslant (4\,r)^{-r}$, $\max (|g_1|, \ldots, |g_r|) = 1$ and $|g_r \, t^r + g_{r+1} \, t^{r+1} + \cdots| \leqslant (4\,r)^{-r}$ on $\mathcal{B}(0,1)$. Then g admits a root on $\mathcal{B}(0,1)$.*

Proof. Let $G(t)=g_1\,t+\cdots+g_r\,t^r$. We may factor $G(t)=(t-\alpha_1)\cdots(t-\alpha_r)$ with $\alpha_1=0$. Let $0<\rho\leqslant 1$ be such that $|\rho-|\alpha_i||\geqslant\frac{1}{2\,r-1}>\frac{1}{r}$ for all i. Then we have $|G(t)|\geqslant(2\,r)^{-r}$ for all $t\in\mathbb{C}$ with $|t|=\rho$, whence $|g(t)-G(t)|\leqslant 2\,(4\,r)^{-r}<(2\,r)^{-r}\leqslant|G(t)|$. By Rouché's theorem, it follows that g and G admit the same number of zeros in $\mathcal{B}(0,r)$. Hence g admits at least one zero inside $\mathcal{B}(0,r)\subseteq\mathcal{B}(0,1)$. □

LEMMA 15. *There exist positive constants ρ', C and ν such that*

$$|f(z)|\leqslant C\,|E_1(z)\,z^{\nu}\,2^{-nr}|\;\Longrightarrow\;\exists z'\in\mathbb{C}, |z'-z|\leqslant 2^{-n}\wedge f(z')=0$$

for all $n\geqslant\lceil r\log_2(4\,r)\rceil$ and $z\in\mathcal{S}_{\mathrm{Id}}$ with $2^{-n}\leqslant|z|\leqslant\rho'$.

Proof. Let C and ν be as in Theorem 11 and ϱ, A and τ as in Lemma 13. Take $\rho'=\min(\rho,\varrho)$. We thus have $\|\Phi_0\|\geqslant C\,|E_1(z)\,z^{\nu}|$ and $|f_{r+k}|\leqslant\|\Phi_0\|\,|A\,z^{\tau}|^{k-r+1}$ for all k. Let $g(t)=f(z+2^{-n}\,t)=g_0+g_1\,t+\cdots$. Then it follows that $M:=\max(|g_0|,\ldots,|g_{r-1}|)\geqslant\|\Phi_0\|\,2^{-(r-1)n}$ and $|g_k|\leqslant\|\Phi_0\|\,|A\,z^{\tau}|^{k-r}\,2^{-kn}\leqslant M\,|A\,z^{\tau}|^{k-r+1}\,2^{-n}$ for $k\geqslant r$. Assuming that $|f(z)|\leqslant C\,|E_1(z)\,z^{\nu}\,2^{-nr}|$, we also obtain $|g_0|\leqslant\|\Phi_0\|\,2^{-nr}\leqslant M\,2^{-n}\leqslant M\,(4\,r)^{-r}$. Decreasing ρ' if necessary, we may arrange ourselves so that $|A\,z^{\tau}|\leqslant 1/2$. Consequently, $\delta(t)=g_r\,t^r+g_{r+1}\,t^{r+1}+\cdots$ satisfies $|\delta(t)|\leqslant M\,2^{-n}\leqslant M\,(4\,r)^{-r}$ for $|t|\leqslant 1$. We now conclude by Lemma 14. □

4.3. Global uniform complexity bounds

We are now in a position to prove our main theorem. We start with proving the uniform bound on "super-admissible" sectors near singularities. Here the sector $\mathcal{S}=\mathcal{S}_{\rho,\theta,\alpha}$ is said to be *super-admissible* if we may take $\rho'=\rho$ in Lemma 15, as well as in the analoguous statement on $\mathcal{S}_{\sigma}$ for each permutation σ of $\{1,\ldots,r\}$. Given $\varepsilon>0$ and $z,z'\in\mathbb{C}$ with $|z'-z|\leqslant\varepsilon$, we will say that z' is an *ε-approximation* of z.

LEMMA 16. *Assume that 0 is a singularity for L and that f is a solution to $L\,f=0$ on a super-admissible sector $\mathcal{S}=\mathcal{S}_{\rho,\theta,\alpha}$, with holonomic initial conditions at a point in $\mathcal{S}\cap\mathbb{K}$. Denote $\Theta=\{u\in\mathcal{S}: f(u)=0\}$. Then there exists an algorithm that takes $n\in\mathbb{N}$ and $z\in\mathcal{S}\cap\mathbb{D}[\mathrm{i}]$ with $\mathcal{B}(z,2^{-n})\cap(\partial\mathcal{S}\cup\Theta)=\varnothing$ and $\mathrm{size}(z)\leqslant n$ on input and that computes $v\in\mathbb{F}[\mathrm{i}]$ on output with $|f(z)-v|\leqslant 2^{-n}\,|f(z)|$. Moreover, the running time of the algorithm is bounded by $O(\mathsf{I}(n)\log^3 n)$, uniformly in z.*

Proof. Let φ_i and b_i denote the accelero-sums of $\tilde{\varphi}_i$ and $\tilde{b}_i$ on $\mathcal{S}$. By Theorem 10, we may compute 2^{-n}-approximations of the evaluations $\varphi_i(z)$ in time $O(\mathsf{I}(n)\log^3 n)$, uniformly for $z\in\mathcal{S}\cap\mathbb{D}[\mathrm{i}]$. In particular, the constants $\lambda_1,\ldots,\lambda_r$ with $f=\lambda_1\,b_1+\cdots+\lambda_r\,b_r$ can be evaluated with a precision of n bits in time $O(\mathsf{I}(n)\log^3 n)$.

For a given $z\in\mathcal{S}\cap\mathbb{D}[\mathrm{i}]$, we first determine a permutation π such that $z\in\mathcal{S}_{\pi}$. Modulo a permutation of the basis elements b_i, we may assume without loss of generality that $\pi=\mathrm{Id}$. In order to evaluate f at z, we perform tentative evaluations at increasing bit precisions $n'=n,2\,n,4\,n,\ldots$ until the desired approximation with a relative precision of n bits is found. For the tentative evaluations, we proceed as follows:

- We compute $2^{-n'-1}$-approximations of $\varphi_1(z),\ldots,\varphi_r(z)$.
- We compute $2^{-n'}$-approximations of $\varphi_2(z)\,E_2(z)/E_1(z),\ldots,\varphi_n(z)\,E_r(z)/E_1(z)$.
- Summing up, we obtain a $r\,2^{-n'}$-approximation of $f(z)/E_1(z)$.

If the $r\ 2^{-n'}$-approximation of $f(z)\,/\,E_1(z)$ has a relative precision of at least $n+1$ bits, then we obtain v using one final multiplication with a floating point approximation of $E_1(z)$. If $f(z)/E_1(z)$ has a smaller relative precision, then we set $n' := 2\,n'$ and keep iterating.

Now whenever both $n' \geqslant r\,n - \lfloor \nu \log_2 |z| + \log_2 C \rfloor$ and $|f(z)\,/\,E_1(z)| \leqslant 2^{-n'}$, Lemma 15 implies that $\mathcal{B}(z, 2^{-n}) \cap \Theta \neq \varnothing$. In other words, the iteration will stop whenever $n' \geqslant r\,n - \lfloor \nu \log_2 |z| + \log_2 C \rfloor + \log_2 r$. Since $|z| \geqslant 2^{-n}$, this happens for $n' = O(n)$. Since we double n' at every iteration, the total running time is dominated by the running time of the last tentative evaluation at precision $n' = O(n)$. The most expensive step of this tentative evaluation is the computation of the $2^{-n'}$-approximations of $\varphi_1(z), \ldots, \varphi_r(z)$. By Theorem 10, this can be done in time $O(\mathsf{I}(n') \log^3 n') = O(\mathsf{I}(n) \log^3 n)$, uniformly in z. □

Proof of Theorem 1. Let $\sigma \in \Sigma$ be one of the singularities and let $\mathcal{S}_1 \cup \cdots \cup \mathcal{S}_\ell$ be an admissible ball cover in the neighbourhood of σ. For each admissible sector $\mathcal{S}_i$ and each connected component $\mathcal{C}$ of $\mathcal{S}_i \cap \Omega$ (there are at most two such connected components), we also arbitrarily pick a point $z_{\mathcal{C}}$ in $\mathcal{C} \cap \mathbb{D}[\mathrm{i}]$. We may compute 2^{-n}-approximations for $f(z_{\mathcal{C}}), \ldots, f^{(r-1)}(z_{\mathcal{C}})$ in time $O(\mathsf{I}(n) \log^2 n)$. These values may be used as initial conditions for f on $\mathcal{S}_i$.

For $z \in \Omega \cap \mathbb{D}[\mathrm{i}]$ sufficiently close to σ, we use the following algorithm for the evaluation of $f(z)$. Among the sectors $\mathcal{S}_i$ that contain z, we pick the one for which $d(z, \partial \mathcal{S}_i)$ is maximal. In particular, $d(z, \partial \mathcal{S}_i) \geqslant \gamma_i\,|z - \sigma|$ for some fixed constant $\gamma_i > 0$. Let $\mathcal{C}$ be the connected component of of $\mathcal{S}_i \cap \Omega$ that contains z. We now evaluate $f(z)$ using the algorithm from Lemma 16, by using the initial conditions for f at $z_{\mathcal{C}}$. Applying Lemma 16 on each of the sectors $\mathcal{S}_i$, we obtain a constant r_σ such that $f(z)$ can be approximated with a relative precision of n bits in time $O(\mathsf{I}(n) \log^3 n)$, uniformly in $z \in \Omega \cap \mathbb{D}[\mathrm{i}] \cap \mathcal{B}(\sigma, r_\sigma)$, provided that $\mathcal{B}(z, 2^{-n}) \cap (\partial \Omega \cup \Theta) = \varnothing$.

Considering the change of variables $z \to 1\,/\,z$, one may prove in a similar way that, for some sufficiently large R, we can approximate $f(z)$ with a relative precision of n bits in time $O(\mathsf{I}(n) \log^3 n)$, uniformly in $z \in \Omega \cap \mathbb{D}[\mathrm{i}] \cap \{u \in \mathbb{C} \colon |u| \geqslant R\}$, provided that $\mathcal{B}(z, 2^{-n}) \cap (\partial \Omega \cup \Theta) = \varnothing$.

Let $U = \{u \in \mathbb{C} \colon |u| > R \wedge (\forall \sigma \in \Sigma, |u - \sigma| < r_\sigma)\}$. The complement $\mathbb{C} \setminus U$ is a compact set that contains none of the singularities of f. Using the complexity bounds from [7], it follows that a 2^{-n}-approximation for $f(z)$ can be computed in time $O(\mathsf{I}(n) \log^2 n)$, uniformly in $z \in (\mathbb{C} \setminus U) \cap \Omega \cap \mathbb{D}[\mathrm{i}]$. Now $f(z)$ admits only a finite number of zeros on $\mathbb{C} \setminus U$ and each zero has multiplicity at most $r - 1$. Considering the local power series expansions around any of these zeros ω, we observe that $|f(z)| > c\,|z - \omega|^r$ for some computable contant $c > 0$ and z sufficiently close to ω. Provided that $\mathcal{B}(z, 2^{-n}) \cap (\partial \Omega \cup \Theta) = \varnothing$, this implies that we can also compute an approximation for $f(z)$ with a relative precision of n bits in time $O(\mathsf{I}(n) \log^2 n)$, uniformly for $z \in (\mathbb{C} \setminus U) \cap \Omega \cap \mathbb{D}[\mathrm{i}]$. □

5. FURTHER THOUGHTS AND CHALLENGES

There are several directions in which the results of this paper can be extended or made more precise.

More general constants. In our main Theorem 1, we assumed that $\mathbb{K}$ is the field of algebraic numbers. Following the Chudnovsky's [3], and using the baby-step-giant-step technique, one may replace $\mathbb{K}$ with more general effective subfield of $\mathbb{C}$ whose elements can be approximated fast. More precisely, if for any constants z in $\mathbb{K}$ we can compute a 2^{-n}-approximation of z in time $O(\mathsf{I}(n^{3/2}) \log^2 n)$, then Theorem 1 still holds, but one should replace the uniform complexity bound $O(\mathsf{I}(n) \log^3 n)$ by $O(\mathsf{I}(n^{3/2}) \log^2 n)$.

Riemann surfaces. In this paper, we used Ω for the domain of our holonomic function f. Of course, f is really defined on the covering space of $\mathbb{C} \setminus \Sigma$ which is a Riemann surface. Points on this Riemann surface can be represented by broken line paths as in [7, 8, 9]. By using a suitable size function for broken line paths with vertices in $\mathbb{D}[\mathrm{i}]$, one may extend Theorem 1 to the evaluation of f at points above $\mathbb{D}[\mathrm{i}]$ on this Riemann surface.

Fast approximation of zeros. Given a sufficiently good approximation $\tilde{z} \in \mathbb{D}[\mathrm{i}]$ of a zero z of f of multiplicity μ (we must have $\mu < r$), we may use Newton method's $\tilde{z}' := \tilde{z} - \mu\, f(\tilde{z}) / f'(\tilde{z})$ to compute a better approximation $\tilde{z}'$. Since the evaluations of f and f' can be done with good uniform complexity, this should make it possible to compute a 2^{-n}-approximation of z in time $O(\mathsf{l}(n) \log^3 n)$, uniformly in $\tilde{z}$ under suitable conditions. It would be a useful contribution to prove a more precise statement of this kind.

Ball evaluations. In this paper, we assumed that the points z where we evaluate f are exactly known. An interesting question concerns the efficient computation of high quality ball lifts $\boldsymbol{f}$ of f. In that case, the evaluation point z is replaced by an explicit ball $\boldsymbol{z} = \mathcal{B}(z, \rho)$ with $z \in \mathbb{D}[\mathrm{i}]$ and $\rho \in \mathbb{D}^{\geqslant}$, and the evaluation $\boldsymbol{f}(\boldsymbol{z})$ should be a similar ball $\boldsymbol{u} = \mathcal{B}(u, \sigma)$ with the property that $f(\boldsymbol{z}) \subseteq \boldsymbol{u}$ and $f(\boldsymbol{z})$ contains two points with distance at least σ. It would be worthwhile to extend Theorem 1 to this kind of arithmetic.

Multi-summation. When introducing the theory of accelero-summability [4, 5], Écalle also described a variant which only relies on the evaluation of iterated Laplace integrals (instead of the more general accelerations). This idea was further developed by Balser [1] who rebaptized it under the term “multi-summation”. It is quite plausible that [9] and the present paper can be adapted to this setting.

BIBLIOGRAPHY

[1] W. Balser. *From divergent power series to analytic functions. Theory and application of multisummable power series*, volume 1582 of *Lect. Notes in Math.* Springer-Verlag, Berlin, 1994.

[2] R. P. Brent. The complexity of multiprecision arithmetic. In *Complexity of Computational problem solving*, 1976.

[3] 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.

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

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

[6] B. Haible and T. Papanikolaou. Fast multiple-precision evaluation of elementary functions. Technical Report TI-7/97, Universität Darmstadt, 1997.

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

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

[9] J. van der Hoeven. Efficient accelero-summation of holonomic functions. *JSC*, 42(4):389–428, 2007.

[10] E. A. Karatsuba. Fast evaluation of Bessel functions. *Integral Transforms and Special Functions*, 1(4):269–276, 1993.

[11] Marc Mezzarobba. *Autour de l'évaluation numérique des fonctions D-finies*. PhD thesis, École polytechnique, 2011.

[12] H. Poincaré. Sur les intégrales irrégulières des équations linéaires. *Acta Math.*, 8:295–344, 1886.