

Optimal contours for high-order derivatives

FOLKMAR BORNEMANN* AND GEORG WECHSLBERGER

Zentrum Mathematik – M3, Technische Universität München, 80290 München, Germany

*Corresponding author: bornemann@tum.de wechslbe@ma.tum.de

[Received on 3 February 2012; revised on 31 July 2012; accepted on 31 July 2012]

As a model of more general contour integration problems we consider the numerical calculation of high-order derivatives of holomorphic functions using Cauchy’s integral formula. Bornemann (2011, Accuracy and stability of computing high-order derivatives of analytic functions by Cauchy integrals. *Found. Comput. Math.*, **11**, 1–63) showed that the condition number of the Cauchy integral strongly depends on the chosen contour and solved the problem of minimizing the condition number for circular contours. In this paper, we minimize the condition number within the class of grid paths of step size h using Provan’s algorithm for finding a shortest enclosing walk in weighted graphs embedded in the plane. Numerical examples show that optimal grid paths yield small condition numbers even in those cases where circular contours are known to be of limited use, such as for functions with branch-cut singularities.

Keywords: high-order derivatives; optimal contours; condition number; Cauchy integral; shortest enclosing walk; Provan’s algorithm.

1. Introduction

To escape from the ill conditioning of difference schemes for the numerical calculation of high-order derivatives, numerical quadrature applied to Cauchy’s integral formula has on various occasions been suggested as a remedy (for a survey of the literature, see Bornemann, 2011). To be specific, we consider a function f that is holomorphic on a complex domain $D \ni 0$; Cauchy’s formula gives¹

$$f^{(n)}(0) = \frac{n!}{2\pi i} \int_{\Gamma} z^{-n-1} f(z) dz \quad (1.1)$$

for each cycle $\Gamma \subset D$ that has winding number $\text{ind}(\Gamma; 0) = 1$. If Γ is not carefully chosen, however, the integrand tends to oscillate at a frequency of order $\mathcal{O}(n^{-1})$ with very large amplitude (Bornemann, 2011, Fig. 4). Hence, in general, there is much cancellation in the evaluation of the integral and ill conditioning returns through the back door. The condition number of the integral² is (Deuffhard & Hohmann, 2003, Lemma 9.1)

$$\kappa(\Gamma, n) = \frac{\int_{\Gamma} |z|^{-n-1} |f(z)| |dz|}{|\int_{\Gamma} z^{-n-1} f(z) dz|}$$

¹ Without loss of generality, we evaluate derivatives at $z=0$.

² Given an accurate and stable (i.e., with positive weights) quadrature method such as Gauss–Legendre or Clenshaw–Curtis, this condition number actually yields, by # loss of significant digits $\approx \log_{10} \kappa(\Gamma, n)$, an estimate of the error caused by round-off in the last significant digit of the data (i.e., the function f).

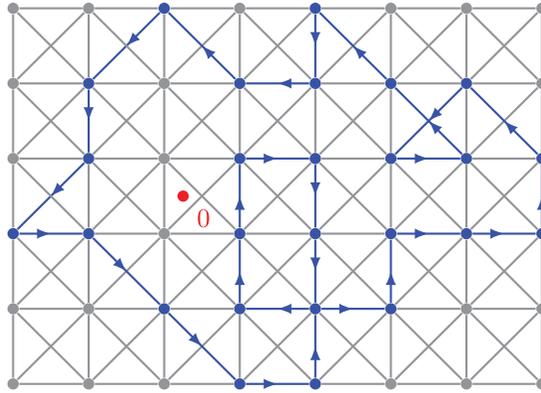


FIG. 1. Path Γ with $\text{ind}(\Gamma; 0) = 1$ in a grid-graph of step size h .

and Γ should be chosen so as to make this number as small as possible. Equivalently, since the denominator is, by Cauchy's theorem, independent of Γ , we have to minimize

$$d(\Gamma) = \int_{\Gamma} |z|^{-n-1} |f(z)| \, d|z|. \quad (1.2)$$

Bornemann (2011) considered circular contours of radius r ; he found that there is a unique $r_* = r(n)$ solving the minimization problem and that there are different scenarios for the corresponding condition number $\kappa_*(n)$ as $n \rightarrow \infty$:

- $\kappa_*(n) \rightarrow \infty$ if f is in the Hardy space H^1 ;
- $\limsup_{n \rightarrow \infty} \kappa_*(n) \leq M$ if f is an entire function of completely regular growth, which satisfies a nonresonance condition of the zeros and whose Phragmén–Lindelöf indicator possesses M maxima (a small integer).

Hence, though those (and similar) results basically solve the problem of choosing proper contours for entire functions, much better contours have to be found for the class H^1 . Moreover, the restriction to circles lacks any algorithmic flavour that would point to more general problems depending on the choice of contours, such as the numerical solution of highly oscillatory Riemann–Hilbert problems (Olver, 2011).³

In this paper, we solve the contour optimization problem within the more general class of grid paths of step size h (see Fig. 1; we allow diagonals to be included) as they are known from Artin's proof of the general, homological version of Cauchy's integral theorem (Lang, 1999, IV.3). Such paths are composed from horizontal, vertical and diagonal edges taken from a (bounded) grid $\Omega_h \subset D$ of step size h . Now, the weight function (1.2), being *additive* on the Abelian group of path chains, turns the grid Ω_h into an edge-weighted graph such that each optimal grid path W_* becomes a *shortest enclosing walk* (SEW); 'enclosing' because we have to match the winding number condition $\text{ind}(W_*; 0) = 1$. An

³ Taking the contour optimization developed in this paper as a model, Wechsberger & Bornemann (2012) have recently addressed the deformation of Riemann–Hilbert problems from an algorithmic point of view.

efficient solution of the SEW problem for embedded graphs was found by Provan (1989) and serves as a starting point for our work.

Outline of the Paper. In Section 2, we discuss general embedded graphs in which an optimal contour is to be searched for; we discuss the problem of finding an SEW and recall Provan's algorithm. In Section 3, we discuss some implementation details and tweaks for the problem at hand. Finally, in Section 4, we give some numerical examples; these can easily be constructed in a way that the new algorithm outperforms, by orders of magnitude, the optimal circles of Bornemann (2011) with respect to the accuracy and direct symbolic differentiation with respect to efficiency.

2. Contour graphs and the SEW

By generalizing the grid Ω_h , we consider a finite graph $G = (V, E)$ embedded in D , that is, built from vertices $V \subset D$ and edges E that are smooth curves connecting the vertices within the domain D . We write uv for the edge connecting the vertices u and v ; by (1.2), its weight is defined as

$$d(uv) = \int_{uv} |z|^{-n-1} |f'(z)| |dz|. \quad (2.1)$$

A walk W in the graph G is a closed path built from a sequence of adjacent edges, written as (where $\dot{+}$ denotes joining of paths)

$$W = v_1 v_2 \dot{+} v_2 v_3 \dot{+} \cdots \dot{+} v_m v_1;$$

it is said to *enclose* the obstacle 0 if the winding number is $\text{ind}(W; 0) = 1$. The set of all possible enclosing walks is denoted by Π . As discussed in Section 1, the condition number is optimized by the (not necessarily unique) SEW

$$W_* = \underset{W \in \Pi}{\text{argmin}} d(W)$$

where, with $W = v_1 v_2 \dot{+} v_2 v_3 \dot{+} \cdots \dot{+} v_m v_1$ and $v_{m+1} = v_1$, the *total weight* is

$$d(W) = \sum_{j=1}^m d(v_j v_{j+1}).$$

The problem of finding such an SEW was solved by Provan (1989): the idea is that with $\mathcal{P}_{u,v}$ denoting a shortest path between u and v , any SEW $W_* = w_1 w_2 \dot{+} w_2 w_3 \dot{+} \cdots \dot{+} w_m w_1$ can be cast in the form (Provan, 1989, Theorem 1)

$$W_* = \mathcal{P}_{w_1, w_j} \dot{+} w_j w_{j+1} \dot{+} \mathcal{P}_{w_{j+1}, w_1}$$

for at least one j . Hence, any SEW W_* is already specified by one of its vertices and one of its edges; therefore

$$W_* \in \tilde{\Pi} = \{\mathcal{P}_{u,v} \dot{+} vw \dot{+} \mathcal{P}_{w,u} : u \in V, vw \in E\}.$$

Provan's algorithm finds W_* by (i) building the finite set $\tilde{\Pi}$; (ii) removing all walks from it that do not enclose $z = 0$; and (iii) selecting a walk from the remaining candidates that has the lowest total weight. Using the Fredman & Tarjan (1987) implementation of Dijkstra's algorithm to compute the shortest paths $\mathcal{P}_{u,v}$, the run time of the algorithm is known to be (Provan, 1989, Corollary 2)

$$\mathcal{O}(|V| |E| + |V|^2 \log |V|). \quad (2.2)$$

3. Implementation details

We restrict ourselves to graphs Ω_h given by finite square grids of step size h , centred at $z = 0$ —with all vertices and edges removed that do not belong to the domain D . Since Provan’s algorithm just requires an embedded graph but not a planar graph, we may add the diagonals of the grid cells as further edges to the graph (see Fig. 1).⁴ For such a graph Ω_h , with or without diagonals, we have $|V| = \mathcal{O}(h^{-2})$ and $|E| = \mathcal{O}(h^{-2})$ so that the complexity bound (2.2) simplifies to

$$\mathcal{O}(h^{-4} \log h^{-1}).$$

3.1 Edge-weight calculation

Using the edge weights $d(uv)$ on Ω_h requires the integral in (2.1) to be approximated. Since not much accuracy is needed here,⁵ a simple trapezoidal rule with two nodes is generally sufficient:

$$\begin{aligned} d(uv) &= \int_{uv} |z|^{-(n+1)} |f(z)| \, d|z| \\ &= \frac{|u-v|}{2} (d(u) + d(v)) + \mathcal{O}(h^3) = \tilde{d}(uv) + \mathcal{O}(h^3) \end{aligned}$$

with the *vertex weight*

$$d(z) = |z|^{-(n+1)} |f(z)|. \quad (3.1)$$

Although $\tilde{d}(uv)$ will typically have an accuracy of not more than just a few bits for the rather coarse grids Ω_h we work with, we have not encountered a single case in which a more accurate computation of the weights would have resulted in a different SEW W_* .

3.2 Reducing the size of $\tilde{\Pi}$

As described in Section 2, Provan’s algorithm starts by building a walk for every pair $(v, e) \in V \times E$ and then proceeds by selecting the best enclosing one. A simple heuristic, which worked well for all our test cases, helps to considerably reduce the number of walks to be processed: let

$$v_* = \operatorname{argmin}_{v \in V} d(v)$$

and define W_{v_*} as an SEW subject to the constraint

$$W_{v_*} \in \tilde{\Pi}_{v_*} = \{\mathcal{P}_{v_*,u} \dot{+} uw \dot{+} \mathcal{P}_{w,v_*} : uw \in E\}.$$

Obviously, W_* and W_{v_*} do not need to agree in general, as v_* does not have to be traversed by W_* . However, since v_* is the vertex with lowest weight, both walks differ mainly in a region that has no, or very minor, influence on the total weight and consequently, also no significant influence on the condition number. Actually, W_* and W_{v_*} yielded precisely the same total weight for all functions that we

⁴ These diagonals increase the number of possible slopes which results, e.g., in improved approximations of the direction of steepest descent at a saddle point of $d(z)$ (Bornemann, 2011, Section 9) or in a faster U-turn around the end of a branch cut; see Fig. 5. The latter case leads to some significant reductions in the condition number; see Fig. 4.

⁵ Recall that optimizing the condition number is just a question of order of magnitude but not of precise numbers. Once the contour Γ has been fixed, a much more accurate quadrature rule will be employed to calculate the integral (1.1) itself; see Section 3.5.

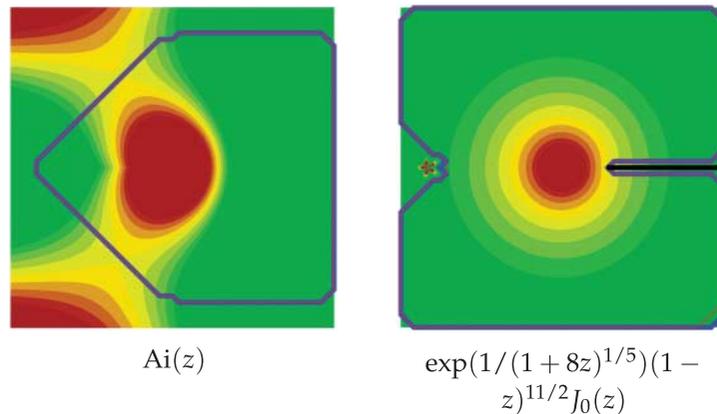


FIG. 2. W_* (red) W_{v^*} (blue): the colour coding shows the size of $\log d(z)$ with red for large values and green for small values. The smallest level shown is the threshold, below which the edges of W_* do not contribute to the first couple of significant digits of the total weight. The plots illustrate that W_* and W_{v^*} differ typically in just a small region well below this threshold; consequently, both walks yield about the same condition number. On the right, note the *five-leaved clover* that represents the combination of algebraic and essential singularity at $z = -1$.

have studied (Fig. 2 compares W_* with W_{v^*} for two typical examples). Using that heuristic, the run time of Provan's algorithm improves to $\mathcal{O}(|E| + |V| \log |V|)$ because its main part reduces to applying Dijkstra's shortest path algorithm just once. In the case of the grid Ω_h , this bound simplifies to

$$\mathcal{O}(h^{-2} \log h^{-1}).$$

3.3 Size of the grid domain

The side length l of the square domain supporting Ω_h has to be chosen large enough to contain an SEW that would approximate an optimal general integration contour. For example, if f is entire, we choose l large enough for this square domain to cover the optimal circular contour: $l > 2r_*$, where r_* is the optimal radius given in Bornemann (2011); a particularly simple choice is $l = 3r_*$. In other cases, we employ a simple search for a suitable value of l by calculating W_* for increasing values of l until $d(W_*)$ does not decrease substantially any more. During this search, the grid will be just rescaled, that is, each grid uses a *fixed* number of vertices; this way only the number of search steps enters as an additional factor in the complexity bound.

3.4 Multilevel refinement of the SEW

Choosing a proper value of h is not straightforward since we would like to balance a good approximation of a generally optimal integration contour with a reasonable amount of computing time. In principle, we would construct a sequence of SEWs for smaller and smaller values of h until the total weight of W_* does not substantially decrease any more. To avoid an undue amount of computational work, we do not refine the grid everywhere but use an adaptive refinement by confining it to a tubular neighbourhood of the currently given SEW W_* (see Fig. 3):

- (1) calculate W_* within an initial grid;
- (2) subdivide each rectangle adjacent to W_* into four rectangles;

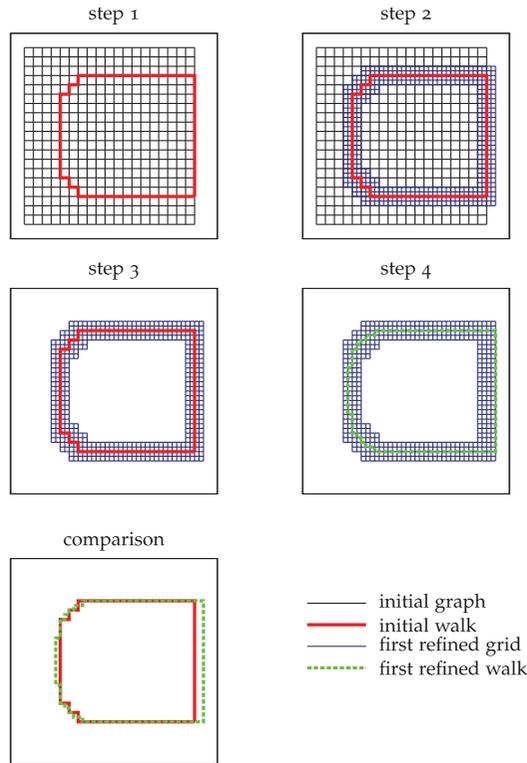


FIG. 3. Multilevel refinement of W_* ($f(z) = 1/\Gamma(z), n = 2006$).

- (3) remove all other rectangles;
- (4) calculate W_* in the newly created graph.

As long as the total weight of W_* decreases substantially, steps 2–4 are repeated. It is even possible to tweak that process further by not subdividing rectangles that containing only vertices or edges of W_* with weights below a certain threshold. By geometric summation, the complexity of the resulting algorithm is

$$\mathcal{O}(H^{-4} \log H^{-1}) + \mathcal{O}(h^{-2} \log h^{-1})$$

where H denotes the step size of the coarsest grid and $h = H/2^k$ the step size after k loops of adaptive refinement. An analogous approach to the constrained W_{v_*} variant of the SEW algorithm given in Section 3.2 reduces the complexity further to

$$\mathcal{O}(H^{-2} \log H^{-1}) + \mathcal{O}(h^{-1} \log h^{-1}),$$

which is close to the best possible bound $\mathcal{O}(h^{-1})$ given by the work that would be needed to just list the SEW.

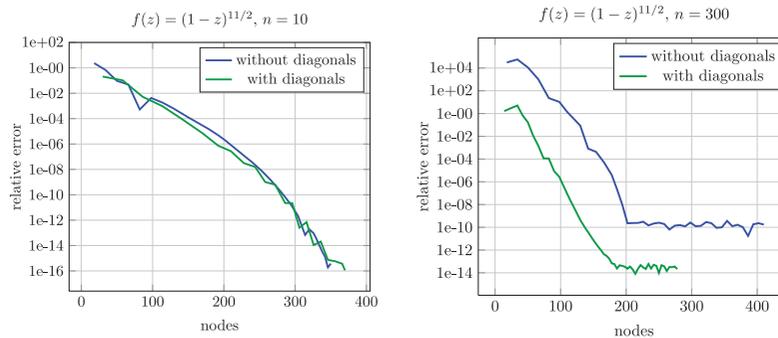


FIG. 4. Illustration of the spectral accuracy of piecewise Clenshaw–Curtis quadrature on SEW contours for a function with a branch-cut singularity. For larger n , we observe a significant improvement by adding diagonals to the grid. We get to machine precision for $n = 10$ and lose about two digits for $n = 300$. (Note that for optimized circular contours the loss would have been about six digits for $n = 10$ and about 15 digits for $n = 300$; cf. Bornemann 2011, Theorem 4.7).

3.5 Quadrature rule for the Cauchy integral

Finally, after calculation of the SEW $\Gamma = W_*$, the Cauchy integral (1.1) has to be evaluated by some *accurate* numerical quadrature. We decompose Γ into maximally straight line segments, each of which can be a collection of many edges. On each of those line segments we employ Clenshaw–Curtis quadrature at Chebyshev–Lobatto points. Additionally, we neglect segments with a weight smaller than 10^{-24} times the maximum weight of an edge of Γ , since such segments will not contribute to the result within machine precision. This way we not only get *spectral accuracy* but also, in many cases, fewer nodes than would be needed by (not truncated) trapezoidal sums on a circular contour: Fig. 4 shows an example with order of differentiation $n = 300$. Here, an accurate solution is obtained with just about 200 nodes — a number, which is well below the sampling condition bound for circular contours (Bornemann, 2011, Section 2.1). Of course, trapezoidal sums would also benefit from some recursive device that helps to neglect those nodes which do not contribute to the numerical result.

4. Numerical results

Table 1 displays condition numbers of SEWs W_* as compared to the optimal circles C_{r_*} for five functions; Table 2 gives the corresponding CPU times and Fig. 5 shows some of the contours. (*All experiments were done using hardware arithmetic.*) The purpose of these examples is twofold, namely, to demonstrate that

- (1) the SEW algorithm *matches* the quality of circular contours in cases where the latter are known to be optimal, such as for entire functions;
- (2) the SEW algorithm is *significantly better* than the circular contours in cases where the latter are known to have severe difficulties.

Thus, the SEW algorithm is a flexible automatic tool that covers various classes of holomorphic functions in a complete algorithmic fashion; in particular, there is no deep theory needed to just let the computation run.

TABLE 1 Condition numbers for some $f(z)$: r_* are the optimal radii given in Bornemann (2011); W_* was calculated in all cases on a 51×51 grid with $l = 3r_*$ (in the last two cases l was found as in Section 3.3)

$f(z)$	n	$\kappa(W_*, n)$	$\kappa(C_{r_*}, n)$
e^z	300	1.1	1.0
$\text{Ai}(z)$	300	1.3	1.2
$1/\Gamma(z)$	300	1.7	1.6
$1/\Gamma(z)$	2006	7.8×10^4	4.7×10^4
$(1-z)^{11/2}$	10	1.4	5.0×10^5
$\exp(1/(1+8z)^{1/5})(1-z)^{11/2}J_0(z)$	100	7.2×10^2	4.3×10^{12}

For $1/\Gamma(z)$, the peculiar order of differentiation $n = 2006$ is one of the very rare resonant cases (specific to this entire function) for which circles give exceptionally large condition numbers (cf. Bornemann, 2011, Table 5). In the last example, differentiation is for $z = 1/\sqrt{2}$.

TABLE 2 CPU times for the examples of Table 1. Here t_{W_*} and $t_{W_{v_*}}$ denote the times to compute W_* and W_{v_*} and t_{quad} denotes the time to approximate the integral (1.1) on such a contour by quadrature (there is no difference between W_* and W_{v_*} ‘point of view of quadrature’?; see Fig. 2.)

$f(z)$	n	Grid	t_{W_*} (s)	$t_{W_{v_*}}$ (s)	t_{quad} (s)
e^z	300	51×51	4.4×10^2	1.5	0.3
$\text{Ai}(z)$	300	25×25	2.1×10^1	0.5	1.7
$\text{Ai}(z)$	300	51×51	4.0×10^2	2.1	2.1
$1/\Gamma(z)$	300	25×25	2.0×10^1	0.5	1.5
$1/\Gamma(z)$	300	51×51	3.6×10^2	2.4	1.3
$1/\Gamma(z)$	2006	51×51	3.6×10^2	2.3	3.1
$(1-z)^{11/2}$	10	51×51	1.4×10^3	5.9	0.2
$\exp(1/(1+8z)^{1/5})(1-z)^{11/2}J_0(z)$	100	51×51	7.0×10^2	3.5	0.3

In the last example, differentiation is for $z = 1/\sqrt{2}$. The timings for the grids of size 25×25 and 51×51 match nicely the $\mathcal{O}(h^{-4} \log h^{-1})$ complexity for W_* and the $\mathcal{O}(h^{-2} \log h^{-1})$ complexity for W_{v_*} .

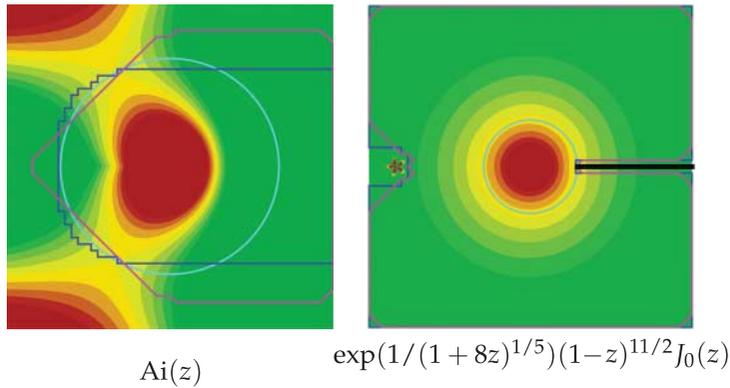


FIG. 5. W_{v_*} (blue: Ω_h without diagonals; magenta: Ω_h with diagonals) vs C_{r_*} (cyan) for some examples from Table 1: the colour coding shows the size of $\log d(z)$ with red for large values and green for small values. The smallest level shown is the threshold below which the edges of W_{v_*} do not contribute to the first significant digits of the total weight.

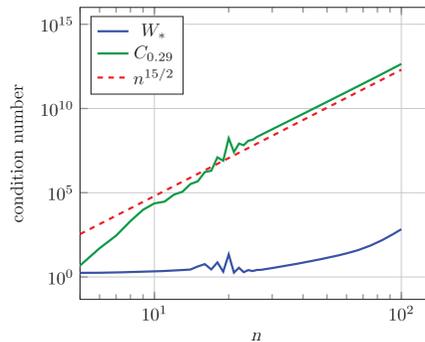


FIG. 6. An example with essential and algebraic singularities: the condition number of the Cauchy integral for $\exp(1/(1+8z)^{1/5})(1-z)^{11/2}J_0(z)$ for varying order of differentiation n at $z=1/\sqrt{2}$; blue: optimal contour W_* in a 51×51 grid graph; green: circular contour with near optimal radius $r=0.29 \approx 1-1/\sqrt{2}$; red: prediction of the growth rate from Bornemann (2011, Theorem 4.7).

In the examples of entire f , we observe that W_* and W_{v_*} , like the optimal circle C_{r_*} would traverse the saddle points of $d(z)$. It was shown in Bornemann (2011, Theorem 10.1) that, for such f , the major contribution of the condition number comes from these saddle points and that circles are (asymptotically, as $n \rightarrow \infty$) paths of steepest descent. Since W_* can cross a saddle point only in a horizontal, vertical or (if enabled) diagonal direction, somewhat larger condition numbers have to be expected. However, the order of magnitude of the condition number of C_{r_*} is precisely matched. This match holds in cases where circles give a condition number of approximately 1, as well as in cases with exceptionally large condition numbers, such as for $f(z) = 1/\Gamma(z)$ in the peculiar case of the order of differentiation $n = 2006$ (cf. Bornemann, 2011, Section 10.4).

For nonentire f , however, optimized circles will be far from optimal in general: Bornemann (2011, Theorem 4.7) shows that the optimized circle C_{r_*} for functions f from the Hardy space H^1 with boundary values in $C^{k,\alpha}$ yields a lower condition-number bound of the form

$$\kappa(C_{r_*}, n) \geq cn^{k+\alpha};$$

for instance, $f(z) = (1-z)^{11/2}$ gives $\kappa(C_{r_*}, n) \sim 0.16059 \cdot n^{13/2}$. On the other hand, W_* gives condition numbers that are orders of magnitude better than those of C_{r_*} by automatically following the branch cut at $(1, \infty)$.

The latter example can easily be cooked up to outperform symbolic differentiation as well: using *Mathematica 8*, the calculation of the n th derivative of $f(z) = \exp(1/(1+8z)^{1/5})(1-z)^{11/2}J_0(z)$ at $z=1/\sqrt{2}$ already takes about a minute for $n=23$ but had to be stopped after *more than a week* for $n=100$. Despite the additional difficulty stemming from the combination of an algebraic and an essential singularity at $z=-1$, the W_{v_*} version of the SEW calculates this $n=100$ derivative to an accuracy of 13 digits in less than 4 s; whereas optimized circular contours would give only about three correct digits here (see Fig. 6).

While many more such numerical experiments would demonstrate that reasonably small condition numbers are obtainable in general,⁶ the study of rigorous condition-number bounds for the SEW has to be postponed to future work.

⁶ The software is provided as a supplement to the e-print version of this paper: arXiv:1107.0498.

REFERENCES

- BORNEMANN, F. (2011) Accuracy and stability of computing high-order derivatives of analytic functions by Cauchy integrals. *Found. Comput. Math.*, **11**, 1–63.
- DEUFLHARD, P. & HOHMANN, A. (2003) *Numerical Analysis in Modern Scientific Computing*, 2nd edn. New York: Springer.
- FREDMAN, M. L. & TARJAN, R. E. (1987) Fibonacci heaps and their uses in improved network optimization algorithms. *J. Assoc. Comput. Mach.*, **34**, 596–615.
- LANG, S. (1999) *Complex Analysis*, 4th edn. New York: Springer.
- OLVER, S. (2011) Numerical solution of Riemann–Hilbert problems: Painlevé II. *Found. Comput. Math.*, **11**, 153–179.
- PROVAN, J. S. (1989) Shortest enclosing walks and cycles in embedded graphs. *Inform. Process. Lett.*, **30**, 119–125.
- WECHSLBERGER, G. & BORNEMANN, F. (2012) Automatic deformation of Riemann-Hilbert problems with applications to the Painlevé II transcendents, arXiv:1206.2446v2.