

A Derivative-Free Algorithm for Computing Zeros of Analytic Functions

P. Kravanja and M. Van Barel, Heverlee

Received November 2, 1998; revised March 30, 1999

Abstract

Let W be a simply connected region in \mathbb{C} , $f : W \rightarrow \mathbb{C}$ analytic in W and γ a positively oriented Jordan curve in W that does not pass through any zero of f . We present an algorithm for computing all the zeros of f that lie in the interior of γ . It proceeds by evaluating certain integrals along γ numerically and is based on the theory of formal orthogonal polynomials. The algorithm requires only f and not its first derivative f' . We have found that it gives accurate approximations for the zeros. Moreover, it is self-starting in the sense that it does not require initial approximations. The algorithm works for simple zeros as well as multiple zeros, although it is unable to compute the multiplicity of a zero explicitly. Numerical examples illustrate the effectiveness of our approach.

AMS Subject Classification: 65H05.

Key Words: Zeros of analytic functions, quadrature method, derivative-free approach, formal orthogonal polynomials.

1. Introduction

Let W be a simply connected region in \mathbb{C} , $f : W \rightarrow \mathbb{C}$ analytic in W and γ a positively oriented Jordan curve in W that does not pass through any zero of f . We consider the problem of computing *all* the zeros of f that lie in the interior of γ . The algorithm that we will present requires only f and not its first derivative f' . It proceeds by evaluating certain integrals along γ numerically and is based on the theory of formal orthogonal polynomials. We have found that our algorithm gives accurate approximations for the zeros. Moreover, it is self-starting in the sense that it does not require initial approximations. The algorithm works for simple zeros as well as multiple zeros, although it is unable to compute the multiplicity of a zero explicitly.

Our approach to the problem of computing all the zeros of an analytic function that lie in the interior of a Jordan curve can be seen as a continuation of the pioneering work of Delves and Lyness [8].

Let N denote the total number of zeros of f that lie in the interior of γ , i.e., the

number of zeros where each zero is counted according to its multiplicity. Suppose from now on that $N > 0$. Delves and Lyness considered the sequence Z_1, \dots, Z_N that consists of all the zeros of f that lie inside γ . Each zero is repeated according to its multiplicity. Suppose that the first derivative f' is available. An easy calculation shows that the logarithmic derivative f'/f has a simple pole at each zero of f with residue equal to the multiplicity of the zero. Cauchy's Theorem implies that

$$N = \frac{1}{2\pi i} \int_{\gamma} \frac{f'(z)}{f(z)} dz. \quad (1)$$

This formula enables one to calculate N via numerical integration. Methods for the determination of zeros of analytic functions that are based on the numerical evaluation of integrals are called *quadrature methods*. A review of such methods was given by Ioakimidis [22]. Delves and Lyness considered the integrals

$$s_p := \frac{1}{2\pi i} \int_{\gamma} z^p \frac{f'(z)}{f(z)} dz, \quad p = 0, 1, 2, \dots$$

The residue theorem implies that the s_p 's are equal to the *Newton sums* of the unknown zeros,

$$s_p = Z_1^p + \dots + Z_N^p, \quad p = 0, 1, 2, \dots \quad (2)$$

These s_p 's can again be calculated via numerical integration along γ .

Delves and Lyness considered the monic polynomial of degree N that has zeros Z_1, \dots, Z_N ,

$$P_N(z) := \prod_{k=1}^N (z - Z_k) =: z^N + \sigma_1 z^{N-1} + \dots + \sigma_N.$$

They called $P_N(z)$ the *associated polynomial* for the interior of γ . Its coefficients can be calculated via Newton's identities.

Theorem (Newton's identities).

$$\begin{aligned} s_1 + \sigma_1 &= 0 \\ s_2 + s_1 \sigma_1 + 2\sigma_2 &= 0 \\ &\vdots \\ s_N + s_{N-1} \sigma_1 + \dots + s_1 \sigma_{N-1} + N \sigma_N &= 0. \end{aligned}$$

Proof: An elegant proof was given by Carpentier and Dos Santos [6]. \square

In this way they reduced the problem to the easier problem of computing the zeros

of a polynomial. Unfortunately, the map from the Newton sums s_1, \dots, s_N to the coefficients $\sigma_1, \dots, \sigma_N$ is usually ill-conditioned. Also, the polynomials that arise in practice may be such that small changes in the coefficients produce much larger changes in some of the zeros. This ill-conditioning of the map between the coefficients of a polynomial and its zeros was investigated by Wilkinson [30]. The location of the zeros determines their sensitivity to perturbations of the coefficients. Multiple zeros and very close zeros are extremely sensitive, but even a succession of moderately close zeros can result in severe ill-conditioning. Wilkinson states that ill-conditioning in polynomials cannot be overcome without, at some stage of the computation, resorting to high precision arithmetic.

If f has many zeros in the interior of γ , then the associated polynomial is of high degree and could be very ill-conditioned. Therefore, if N is large, one has to calculate the coefficients $\sigma_1, \dots, \sigma_N$, and thus the integrals s_1, \dots, s_N , very accurately. To avoid the use of high precision arithmetic and to reduce the number of integrand evaluations needed to approximate the s_p 's, Delves and Lyness suggested to construct and solve the associated polynomial only if its degree is smaller than or equal to a preassigned number M . Otherwise, the interior of γ is subdivided or covered with a finite covering and the smaller regions are treated in turn. The choice of M involves a trade-off. If M is increased, then fewer regions have to be scanned. However, if M is chosen too large, then the resulting associated polynomial may be ill-conditioned. Delves and Lyness chose $M = 5$.

Botten, Craig and McPhedran [2] made a Fortran 77 implementation of the method of Delves and Lyness.

Instead of using Newton's identities to construct the associated polynomial, Li [26] considered (2) as a system of polynomial equations. He used a homotopy continuation method to solve this system.

Recently, the authors (in collaboration with Tetsuya Sakurai) argued that what is wrong with these approaches is that they consider the wrong set of unknowns, cf. our paper [24]. One should consider the mutually distinct zeros and their respective multiplicities *separately*. The quadrature method that we presented is a generalization of the method of Delves and Lyness. It is again based on the numerical evaluation of integrals along γ that involve the logarithmic derivative f'/f , but by using the theory of formal orthogonal polynomials we were able to obtain more accurate approximations for the zeros. The Fortran 90 package ZEAL [25] contains an implementation of our algorithm. In this approach, the mutually distinct zeros are computed by solving a generalized eigenvalue problem whereas the multiplicities are calculated by solving a linear system of equations that has Vandermonde structure. The number of mutually distinct zeros is determined indirectly.

In some applications, the calculation of the derivative f' is more time-consuming than that of f . Delves and Lyness used an integration by parts to derive a formula for s_p that depends only on a multi-valued logarithm of f and not on f' . To

apply this formula, they had to keep track of the sheet on which $\log f(z)$ lies as z runs along the curve γ . Unfortunately, in most cases it is impossible to do this in a completely reliable way, i.e., without accidentally overlooking any sheets. Carpentier and Dos Santos [6] and Davies [7] derived similar formulae. See also Ioakimidis and Anastasselou [23].

The approach that we took in [24] involves integrals along γ that contain the logarithmic derivative f'/f . In this paper we will show that essentially the same results can be obtained via integrals that contain $1/f$. The derivative f' is no longer needed. Of course, in this new approach not the mutually distinct zeros but rather the unknowns Z_1, \dots, Z_N are calculated and the multiplicities cannot be computed explicitly. But apart from this, the algorithm has the same advantages as the algorithm in [24]. In particular, it does not require initial approximations for the zeros and we have found that it gives accurate results.

Formula (1) involves the derivative of f and hence it is no longer available to us for computing the total number of zeros N . Instead one can use the well-known principle of the argument, which states that N is equal to the winding number of the curve $f(\gamma)$ with respect to the origin. This will be discussed in Section 2 along with various reliability issues. In Section 3 we will give an overview of the algorithm for computing zeros of analytic functions that we proposed in [24]. This should enable the reader to contrast this approach with the derivative-free algorithm that we will present in Section 4. We conclude with numerical examples in Section 5.

2. Computing the Total Number of Zeros

The total number of zeros of f that lie inside γ is given by the integral

$$N = \frac{1}{2\pi i} \int_{\gamma} \frac{f'(z)}{f(z)} dz, \quad (3)$$

cf. Eq. (1). Our derivative-free algorithm requires the value of N or an upper bound for it. However, as by assumption the derivative f' is not available to us, we cannot obtain the value of N by evaluating the integral in the right-hand side of (3) numerically. Instead we proceed as follows. By making the substitution $w := f(z)$ we obtain that

$$N = \frac{1}{2\pi i} \int_{f(\gamma)} \frac{1}{w} dw. \quad (4)$$

Here $f(\gamma)$ denotes the image of the curve γ under f . This is a closed curve that avoids the origin. The winding number of $f(\gamma)$ with respect to the origin is defined

as the increase in the argument of $f(z)$ along γ divided by 2π ,

$$n(f(\gamma), 0) := \frac{1}{2\pi} [\arg f(z)]_{z \in \gamma}.$$

Informally speaking, one can say that it is equal to the number of times that the curve $f(\gamma)$ “winds” itself around the origin. A classical theorem in complex analysis (see, e.g., Henrici [20, p. 233]) says that this winding number can be expressed as the integral that appears in the right-hand side of (4). Hence $N = n(f(\gamma), 0)$. This result is known as the “principle of the argument.”

The value of N can thus be computed via an algorithm for computing winding numbers. The range of the function \arg is $(-\pi, \pi]$. If the increase in argument along the straight section

$$[\alpha, \beta] := \{z \in \mathbb{C} : z = t\alpha + (1-t)\beta, \quad 0 \leq t \leq 1\}, \quad \alpha, \beta \in \mathbb{C},$$

satisfies

$$|[\arg f(z)]_{z \in [\alpha, \beta]}| \leq \pi,$$

then

$$[\arg f(z)]_{z \in [\alpha, \beta]} = \arg\left(\frac{f(\beta)}{f(\alpha)}\right),$$

as the reader may easily verify. Let us discretize the curve γ into the sequence of points c_1, \dots, c_G . Define $c_{G+1} := c_1$. Then it follows that

$$N = \frac{1}{2\pi} \sum_{k=1}^G \arg\left(\frac{f(c_{k+1})}{f(c_k)}\right)$$

if

$$|[\arg f(z)]_{z \in [c_k, c_{k+1}]}| \leq \pi \tag{5}$$

for $k = 1, \dots, G$. In other words, if condition (5) is satisfied, then N can be computed simply by evaluating f at the points c_1, \dots, c_G . These considerations form the basis for Henrici’s algorithm [20, pp. 239–241]. See also Ying and Katz [32].

Unfortunately, condition (5) may not be easy to verify for an arbitrary analytic function f . If the discretization of γ is inadequate, then the computed value of $n(f(\gamma), 0)$ and hence N may be wrong. In this sense Henrici’s algorithm is unreliable. Indeed, a finite number of functional or derivative values are not enough to determine the number of zeros of f , even if f is a polynomial. This was shown by Ying in his PhD thesis [31].

There exist several reliable approaches for computing N . However, they all assume that some kind of *global* information is available, which may not always be the case in practice.

- Ying and Katz [32] developed a reliable variant of Henrici's algorithm. They assume that an upper bound for $|f''(z)|$ along an arbitrary line segment is available.
- Herlocker and Ely [21] experimented with a numerical integration approach based on Simpson's rule and the corresponding formula for the integration error. This formula involves the fourth derivative of the integrand evaluated at an unknown point in the integration interval. Automatic differentiation combined with interval arithmetic enabled them to bound the integration error.
- The total number of zeros can also be computed as the topological degree of the mapping

$$F(x, y) := (\operatorname{Re} f(x + iy), \operatorname{Im} f(x + iy))$$

with respect to the interior of γ (interpreted as a subset of \mathbb{R}^2) and the point $(0, 0)$. (We will not go into the details of degree theory. For an excellent introduction, we refer the interested reader to Lloyd's book [27].) Boulton and Sikorski [3] considered the case that γ is the boundary of the unit square $[0, 1] \times [0, 1]$. They proved that, in case F satisfies the Lipschitz condition with constant $K > 0$ and if the infinity norm of F on γ is at least $d > 0$ where $K/(4d) \geq 1$, then at least $4\lfloor K/(4d) \rfloor$ function evaluations are needed to compute the topological degree. See also Traub, Wasilkowski and Woźniakowski [28, pp. 193–194].

These algorithms are reliable but they can only be used if certain global information (an upper bound for the modulus of a higher derivative of f or the Lipschitz constant of $F = (\operatorname{Re} f, \operatorname{Im} f)$) or a lower bound for the infinity norm of F on γ is available or can easily be computed (for example, via automatic differentiation and/or interval arithmetic). If this is not the case and f' is not available, then Henrici's algorithm is really the only algorithm that one can use to compute N . The reader should realize, though, that the computed integer may in fact only be a strictly lower bound for N and hence the output of Henrici's algorithm should be handled with caution. When given an upper bound for N , our algorithm will determine the value of N via the stopping criterion that we will discuss below. However, our algorithm will fail in case the value of N that it is given is in fact strictly smaller than the actual value of N .

An Algorithm Based on f'/f

Let n denote the number of mutually distinct zeros of f that lie inside γ . Let z_1, \dots, z_n be these zeros and v_1, \dots, v_n their respective multiplicities. The quadrature method that we recently proposed in [24] generalizes the approach of Delves and Lyness. Our approach assumes that f as well as its first derivative f' are available. By using the theory of formal orthogonal polynomials, we showed how the mutually distinct zeros can be calculated by solving generalized eigenvalue prob-

lems. The value of n is determined indirectly. Once n and z_1, \dots, z_n have been found, the problem becomes linear and the multiplicities ν_1, \dots, ν_n are computed by solving a linear system of equations that has Vandermonde structure. In this section we will give a brief summary of these results. This should enable the reader to compare this approach with the derivative-free algorithm that we will present in Section 4. For more details (including proofs and a pseudo-code formulation of the algorithm), we refer to [24].

Let \mathbb{P} be the linear space of polynomials with complex coefficients. One defines a symmetric bilinear form

$$\langle \cdot, \cdot \rangle : \mathbb{P} \times \mathbb{P} \rightarrow \mathbb{C}$$

by setting

$$\langle \phi, \psi \rangle := \frac{1}{2\pi i} \int_{\gamma} \phi(z)\psi(z) \frac{f'(z)}{f(z)} dz = \sum_{k=1}^n \nu_k \phi(z_k)\psi(z_k) \quad (6)$$

for any two polynomials $\phi, \psi \in \mathbb{P}$. The latter equality follows from the fact that f'/f has a simple pole at z_k with residue ν_k for $k = 1, \dots, n$. Note that $\langle \cdot, \cdot \rangle$ can be evaluated via numerical integration along γ . In what follows, we will assume that all the “inner products” $\langle \phi, \psi \rangle$ that are needed have been calculated. Let $s_p := \langle 1, z^p \rangle$ for $p = 0, 1, 2, \dots$. These ordinary moments are equal to the *Newton sums* of the unknown zeros,

$$s_p = \sum_{k=1}^n \nu_k z_k^p, \quad p = 0, 1, 2, \dots$$

In particular, $s_0 = \nu_1 + \dots + \nu_n = N$, the total number of zeros. Hence, we may assume that the value of N is known. Let H_k be the $k \times k$ Hankel matrix

$$H_k := \left[s_{p+q} \right]_{p,q=0}^{k-1} = \begin{bmatrix} s_0 & s_1 & \cdots & s_{k-1} \\ s_1 & & \ddots & \vdots \\ \vdots & \ddots & & \vdots \\ s_{k-1} & \cdots & \cdots & s_{2k-2} \end{bmatrix}$$

for $k = 1, 2, \dots$. A monic polynomial φ_t of degree $t \geq 0$ that satisfies

$$\langle z^k, \varphi_t(z) \rangle = 0, \quad k = 0, 1, \dots, t-1, \quad (7)$$

is called a *formal orthogonal polynomial* (FOP) of degree t . (Observe that condition (7) is void for $t = 0$.) The adjective *formal* emphasizes the fact that, in general, the form $\langle \cdot, \cdot \rangle$ does not define a true inner product. An important consequence of this fact is that, in contrast to polynomials that are orthogonal with respect to a true inner product, formal orthogonal polynomials need not exist or need not be unique for every degree. (For details, see Draux [9, 10], Gutknecht [18, 19] or Gragg and

Gutknecht [17].) If (7) is satisfied and φ_t is unique, then φ_t is called a *regular FOP* and t a *regular index*. If we set

$$\varphi_t(z) =: u_{0,t} + u_{1,t}z + \cdots + u_{t-1,t}z^{t-1} + z^t$$

then condition (7) translates into the *Yule–Walker* system

$$\begin{bmatrix} s_0 & s_1 & \cdots & s_{t-1} \\ s_1 & & \ddots & \vdots \\ \vdots & \ddots & & \vdots \\ s_{t-1} & \cdots & \cdots & s_{2t-2} \end{bmatrix} \begin{bmatrix} u_{0,t} \\ u_{1,t} \\ \vdots \\ u_{t-1,t} \end{bmatrix} = - \begin{bmatrix} s_t \\ s_{t+1} \\ \vdots \\ s_{2t-1} \end{bmatrix}. \quad (8)$$

Hence, the regular FOP of degree $t \geq 1$ exists if and only if the matrix H_t is nonsingular.

The following theorem characterizes n , the number of mutually distinct zeros. It enables one, theoretically at least, to calculate n as $\text{rank } H_N$.

Theorem 2. $n = \text{rank } H_{n+p}$ for every nonnegative integer p . In particular, $n = \text{rank } H_N$.

Therefore H_n is nonsingular whereas H_t is singular for $t > n$. Note that $H_1 = [s_0]$ is nonsingular by assumption. The regular FOP of degree 1 exists and is given by $\varphi_1(z) = z - \mu$ where

$$\mu := \frac{s_1}{s_0} = \frac{\sum_{k=1}^n v_k z_k}{\sum_{k=1}^n v_k}$$

is the arithmetic mean of the zeros. Theorem 2 implies that the regular FOP φ_n of degree n exists and tells us also that regular FOPs of degree larger than n do not exist. The polynomial φ_n is easily seen to be

$$\varphi_n(z) = (z - z_1) \cdots (z - z_n). \quad (9)$$

It is the monic polynomial of degree n that has z_1, \dots, z_n as simple zeros. This polynomial has the peculiar property that it is orthogonal to *all* polynomials (including itself),

$$\langle z^p, \varphi_n(z) \rangle = 0, \quad p = 0, 1, 2, \dots \quad (10)$$

Once n is known, the mutually distinct zeros z_1, \dots, z_n can be calculated by solving a generalized eigenvalue problem. Indeed, let $H_n^<$ be the Hankel matrix

$$H_n^< := \begin{bmatrix} s_1 & s_2 & \cdots & s_n \\ s_2 & & \ddots & \vdots \\ \vdots & \ddots & & \vdots \\ s_n & \cdots & \cdots & s_{2n-1} \end{bmatrix}.$$

Theorem 3. *The eigenvalues of the pencil $H_n^< - \lambda H_n$ are given by z_1, \dots, z_n .*

Proof: Let V_n be the Vandermonde matrix based on the points z_1, \dots, z_n ,

$$V_n := \begin{bmatrix} 1 & z_1 & \cdots & z_1^{n-1} \\ \vdots & \vdots & & \vdots \\ 1 & z_n & \cdots & z_n^{n-1} \end{bmatrix},$$

and define the diagonal matrices D_n and $D_n^{(1)}$ as

$$D_n := \text{diag}(v_1, \dots, v_n) \quad \text{and} \quad D_n^{(1)} := \text{diag}(v_1 z_1, \dots, v_n z_n).$$

Then V_n as well as D_n are nonsingular. One can easily verify that the Hankel matrices H_n and $H_n^<$ can be factorized as

$$H_n^< = V_n^T D_n^{(1)} V_n \quad \text{and} \quad H_n = V_n^T D_n V_n.$$

Now let λ^* be an eigenvalue of the pencil $H_n^< - \lambda H_n$ with eigenvector x . Then

$$\begin{aligned} H_n^< x &= \lambda^* H_n x \\ \Leftrightarrow V_n^T D_n^{(1)} V_n x &= \lambda^* V_n^T D_n V_n x \\ \Leftrightarrow D_n^{(1)} y &= \lambda^* D_n y \quad \text{if } y := V_n x \\ \Leftrightarrow \text{diag}(z_1, \dots, z_n) y &= \lambda^* y. \end{aligned}$$

This proves the theorem. \square

Once z_1, \dots, z_n have been found, the multiplicities v_1, \dots, v_n can be computed by solving the Vandermonde system

$$\begin{bmatrix} 1 & \cdots & 1 \\ z_1 & \cdots & z_n \\ \vdots & & \vdots \\ z_1^{n-1} & \cdots & z_n^{n-1} \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \\ \vdots \\ v_n \end{bmatrix} = \begin{bmatrix} s_0 \\ s_1 \\ \vdots \\ s_{n-1} \end{bmatrix}. \tag{11}$$

Note. Vandermonde matrices are often very ill-conditioned [15, 16]. In this case, however, the components of the solution vector of (11) are known to be integers, and therefore there is no problem, even if the linear system (11) happens to be ill-conditioned, as long as the computed approximations for the components of the solution vector have an absolute error that is less than 0.5.

Theorems 2 and 3 suggest the following approach to compute n and z_1, \dots, z_n . Start by computing the total number of zeros N . Next, compute s_1, \dots, s_{2N-2} . As already mentioned, this can be done via numerical integration along γ . The number

of mutually distinct zeros is then calculated as the rank of H_N , $n = \text{rank } H_N$. Finally, the zeros z_1, \dots, z_n are obtained by solving a generalized eigenvalue problem. However, this approach has several disadvantages:

- Theoretically the $N - n$ smallest singular values of H_N are equal to zero. In practice, this will not be the case, and it may be difficult to determine the rank of H_N and hence the value of n in case the gap between the computed approximations for the zero singular values and the nonzero singular values is too small.
- The approximations for z_1, \dots, z_n obtained via Theorem 3 may not be very accurate. Indeed, the mapping from the Newton sums to the zeros and their respective multiplicities,

$$(s_0, s_1, \dots, s_{2n-1}) \mapsto (z_1, \dots, z_n, \nu_1, \dots, \nu_n), \quad (12)$$

is usually very ill-conditioned. (See, e.g., the papers by Gautschi [12, 13, 14] who studied the conditioning of (12) in the context of Gauss quadrature formulae.) Indeed, a classical adage in numerical analysis says that one should avoid the use of ordinary moments.

The algorithm that we proposed in [24] gives more accurate approximations for z_1, \dots, z_n . The idea is the following. The inner products that appear in the Hankel matrices H_n and $H_n^<$ are related to the standard monomial basis. Why not consider a different basis? In other words, why not try to use modified moments instead of ordinary moments? The fact that

$$H_n = \left[\langle z^p, z^q \rangle \right]_{p,q=0}^{n-1} \quad \text{and} \quad H_n^< = \left[\langle z^p, z z^q \rangle \right]_{p,q=0}^{n-1}$$

suggests that one should consider the matrices

$$\left[\langle \psi_p, \psi_q \rangle \right]_{p,q=0}^{n-1} \quad \text{and} \quad \left[\langle \psi_p, \psi_1 \psi_q \rangle \right]_{p,q=0}^{n-1} \quad (13)$$

where ψ_k is a polynomial of degree k for $k = 0, 1, \dots, n - 1$. Of course, even if one succeeds in writing Theorem 3 in terms of the matrices that appear in (13), the question remains which polynomials ψ_k to choose. We found that very accurate results are obtained if one uses the formal orthogonal polynomials. In other words, the zeros of $\varphi_n(z)$ will be computed from inner products that involve $\varphi_0(z), \varphi_1(z), \dots, \varphi_{n-1}(z)$. The value of n will be determined indirectly. Before we can explain this in more detail, we have to say a few words about the orthogonality properties of FOPs.

If H_n is strongly nonsingular, i.e., if all its leading principal submatrices are nonsingular, then we have a full set $\{\varphi_0, \varphi_1, \dots, \varphi_n\}$ of regular FOPs.

What happens if H_n is not strongly nonsingular? By filling up the gaps in the sequence of existing regular FOPs it is possible to define a sequence $\{\varphi_t\}_{t=0}^{\infty}$, with φ_t a monic polynomial of degree t , such that if these polynomials are grouped into blocks according to the sequence of regular indices, then polynomials belonging to different blocks are orthogonal with respect to $\langle \cdot, \cdot \rangle$. More precisely, define $\{\varphi_t\}_{t=0}^{\infty}$ as follows. If t is a regular index, then let φ_t be the regular FOP of degree t . Else define φ_t as $\varphi_r \psi_{t,r}$ where r is the largest regular index less than t and $\psi_{t,r}$ is an arbitrary monic polynomial of degree $t - r$. In the latter case φ_t is called an *inner polynomial*. These polynomials $\{\varphi_t\}_{t=0}^{\infty}$ can be grouped into blocks. Each block starts with a regular FOP and the remaining polynomials are inner polynomials. Note that the last block has infinite length. The block orthogonality property is then expressed by the fact that the *Gram matrix* $G_n := [\langle \varphi_r, \varphi_s \rangle]_{r,s=0}^{n-1}$ is block diagonal. The diagonal blocks are nonsingular, symmetric and zero above the main antidiagonal. (See Bultheel and Van Barel [4] for more details.)

Theorem 3 can be interpreted in the following way: the zeros of the regular FOP of degree n can be calculated by solving a generalized eigenvalue problem. The following theorem shows that this zero/eigenvalue property holds for all regular FOPs. This will enable us to compute regular FOPs in their product representation. The theorem also provides a solution to the problem of how to switch from ordinary moments to modified moments. Define the matrices G_k and $G_k^{(1)}$ as

$$G_k := \left[\langle \varphi_r, \varphi_s \rangle \right]_{r,s=0}^{k-1} \quad \text{and} \quad G_k^{(1)} := \left[\langle \varphi_r, \varphi_1 \varphi_s \rangle \right]_{r,s=0}^{k-1}$$

for $k = 1, 2, \dots$

Theorem 4. *Let $t \geq 1$ be a regular index and let $z_{t,1}, \dots, z_{t,t}$ be the zeros of the regular FOP φ_t . Then the eigenvalues of the pencil $G_t^{(1)} - \lambda G_t$ are given by $\varphi_1(z_{t,1}), \dots, \varphi_1(z_{t,t})$. In other words, they are given by $z_{t,1} - \mu, \dots, z_{t,t} - \mu$ where $\mu = s_1/s_0$.*

Corollary 5. *The eigenvalues of $G_n^{(1)} - \lambda G_n$ are given by $z_1 - \mu, \dots, z_n - \mu$ where $\mu = s_1/s_0$.*

Regular FOPs are characterized by the fact that the determinant of a Hankel matrix is different from zero, while inner polynomials correspond to singular Hankel matrices. To decide whether $\varphi_t(z)$ should be defined as a regular FOP or as an inner polynomial, one could therefore calculate the determinant of H_t and check if it is equal to zero. However, from a numerical point of view such a test “is equal to zero” does not make sense. Because of rounding errors (both in the evaluation of $\langle \cdot, \cdot \rangle$ and in the calculation of the determinant) one would encounter only regular FOPs. Strictly speaking one could say that inner polynomials are not needed in numerical calculations. However, the opposite is true! Let us agree to call a regular FOP *well-conditioned* if its corresponding Yule–Walker system

(8) is well-conditioned, and *ill-conditioned* otherwise. To obtain a numerically stable algorithm, it is crucial to generate only well-conditioned regular FOPs and to replace ill-conditioned regular FOPs by inner polynomials. Stable look-ahead solvers for linear systems of equations that have Hankel structure are based on this principle [1, 5, 11]. In this approach the diagonal blocks in G_n are taken (slightly) larger than strictly necessary to avoid ill-conditioned blocks.

The algorithm for calculating the mutually distinct zeros z_1, \dots, z_n that we proposed in [24] proceeds by computing the polynomials

$$\varphi_0(z), \varphi_1(z), \dots, \varphi_n(z)$$

in their product representation, starting with $\varphi_0(z) \leftarrow 1$ and $\varphi_1(z) \leftarrow z - \mu$. At each step, to decide whether it is numerically feasible to generate the next polynomial in the sequence as a regular FOP, the algorithm uses a heuristic method. By doing a large number of numerical experiments, we reached the conclusion that this heuristic approach leads to accurate results. For more details, we refer to [24].

How does one obtain the value of n ? Theorem 2 and Eqs. (6) and (9) imply the following.

Theorem 6. *Let $t \geq n$. Then $\varphi_t(z_k) = 0$ for $k = 1, \dots, n$ and $\langle z^p, \varphi_t(z) \rangle = 0$ for all $p \geq 0$.*

The value of n can be determined as follows. Suppose that the algorithm has just generated a (well-conditioned) regular FOP $\varphi_r(z)$. To check whether $n = r$, the algorithm scans the sequence

$$\left(|\langle z^\tau \varphi_r(z), \varphi_r(z) \rangle| \right)_{\tau=0}^{N-1-r}. \quad (14)$$

If all the elements are “sufficiently small,” then the algorithm concludes that indeed $n = r$ and it stops.

As we have already mentioned, once n and (approximations for) z_1, \dots, z_n have been found, the multiplicities ν_1, \dots, ν_n are computed by solving the Vandermonde system (11).

This concludes our discussion of the algorithm presented in [24]. From now on we will refer to this algorithm as “algorithm ZEROS/FDF.”

4. A Derivative-Free Algorithm

The results obtained in the previous section are based on the symmetric bilinear form (6), which involves the logarithmic derivative f'/f . Instead, let us consider

the form

$$\langle \cdot, \cdot \rangle_\star : \mathbb{P} \times \mathbb{P} \rightarrow \mathbb{C}$$

defined as

$$\langle \phi, \psi \rangle_\star := \frac{1}{2\pi i} \int_\gamma \phi(z) \psi(z) \frac{1}{f(z)} dz \quad (15)$$

for any two polynomials $\phi, \psi \in \mathbb{P}$. Again, this form can be evaluated via numerical integration along γ and in what follows we will assume that all the “inner products” $\langle \phi, \psi \rangle_\star$ that are needed have been calculated.

We will show that essentially the same results can be obtained with the form $\langle \cdot, \cdot \rangle_\star$ as previously with $\langle \cdot, \cdot \rangle$.

The integrand that appears in the right-hand side of (15) has a pole at every zero of f that lies in the interior of γ and the order of the pole is equal to the multiplicity of the zero. Therefore, the residue theorem implies that $\langle \phi, \psi \rangle_\star$ is equal to the sum of the residues of the function $\phi\psi/f$ at these poles. The following result can easily be verified.

Proposition 7. *If all the N zeros Z_1, \dots, Z_N of f that lie inside γ are simple, then*

$$\langle \phi, \psi \rangle_\star = \sum_{k=1}^N \frac{\phi(Z_k) \psi(Z_k)}{f'(Z_k)}.$$

In general, if f has multiple zeros, then an elegant expression for $\langle \phi, \psi \rangle_\star$ written as a sum is much more difficult to obtain. Fortunately, it is not necessary to have such an expression available. The proofs of Theorems 2 and 3 given in [24] depend completely on the details of the way in which $\langle \phi, \psi \rangle$ can be written as a sum, cf. Eq. (6). However, as we will see, the corresponding theorems can be proved in a different way.

Define $s_p^\star := \langle 1, z^p \rangle_\star$ for $p = 0, 1, 2, \dots$ and let H_k^\star be the $k \times k$ Hankel matrix

$$H_k^\star := \left[s_{p+q}^\star \right]_{p,q=0}^{k-1}$$

for $k = 1, 2, \dots$. The formal orthogonal polynomials associated with $\langle \cdot, \cdot \rangle_\star$ can be defined as before. The coefficients of regular FOPs can be computed by solving a Yule–Walker system, cf. Equation (8). Also, $t \geq 1$ is a regular index if and only if the matrix H_t^\star is nonsingular.

The residue theorem immediately implies that the polynomial

$$P_N(z) = \prod_{k=1}^N (z - Z_k)$$

satisfies

$$\langle z^p, P_N(z) \rangle_\star = 0, \quad p = 0, 1, 2, \dots \quad (16)$$

In this sense, the polynomial $P_N(z)$ behaves with respect to the form $\langle \cdot, \cdot \rangle_\star$ in the same way as the polynomial $\varphi_n(z)$ behaves with respect to $\langle \cdot, \cdot \rangle$, cf. Eqs. (9) and (10). We will prove that N is the largest regular index for $\langle \cdot, \cdot \rangle_\star$. This will enable us to compute the zeros of the regular FOP $P_N(z)$, i.e., the zeros Z_1, \dots, Z_N , in essentially the same way as algorithm ZEROS/FDF computes the zeros of $\varphi_n(z)$, i.e., the mutually distinct zeros z_1, \dots, z_n .

The following lemma will play an important role. Define the set \mathcal{I} as follows:

$$\mathcal{I} := \{ \phi \in \mathbb{P} : \langle z^p, \phi(z) \rangle_\star = 0 \text{ for } p = 0, 1, 2, \dots \}.$$

Lemma 8. *The set \mathcal{I} is equal to the ideal generated by the polynomial P_N . In other words,*

$$\mathcal{I} = \{ \phi \in \mathbb{P} : \exists \alpha \in \mathbb{P} : \phi = \alpha P_N \}.$$

Proof: Suppose that $a \in \mathbb{C}$ lies in the interior of γ . Let the function $g : W \rightarrow \mathbb{C}$ be meromorphic and suppose that g has neither zeros nor poles on γ . Then the coefficient of $(z - a)^{-p-1}$ in the Laurent expansion of g at the point a is given by the integral

$$\frac{1}{2\pi i} \int_\gamma (z - a)^p g(z) dz$$

for $p = 0, 1, 2, \dots$. Let $\phi \in \mathcal{I}$. Then

$$\langle (z - Z_k)^p, \phi(z) \rangle_\star = \frac{1}{2\pi i} \int_\gamma (z - Z_k)^p \frac{\phi(z)}{f(z)} dz = 0$$

for $k = 1, \dots, N$ and $p = 0, 1, 2, \dots$ and thus the function ϕ/f has a removable singularity at the points Z_1, \dots, Z_N . Thus ϕ has to be a multiple of P_N . This proves the lemma. \square

Theorem 9. *The matrix H_N^\star is nonsingular.*

Proof: We will prove that P_N is the only monic polynomial of degree N that is orthogonal to all polynomials of lower degree. Suppose that Q_N is another such polynomial. Then $P_N - Q_N$ is of degree at most $N - 1$ and hence $\langle P_N - Q_N, Q_N \rangle_\star = 0$. Equation (16) then implies that $\langle Q_N, Q_N \rangle_\star = 0$. Thus Q_N is not only orthogonal to all polynomials of degree $\leq N - 1$ but also to all polynomials of degree N . The polynomial $zP_N - zQ_N$ has degree $\leq N$ and therefore $\langle zP_N - zQ_N, Q_N \rangle_\star = 0$. As $\langle zP_N, Q_N \rangle_\star = \langle P_N, zQ_N \rangle_\star = 0$, it follows that $\langle zQ_N, Q_N \rangle_\star = 0$. Thus Q_N is also orthogonal to all polynomials of degree $N + 1$. By continuing this way, one can prove that Q_N is orthogonal to *all* polynomials, $Q_N \in \mathcal{I}$. As Q_N is a

monic polynomial of degree N , the previous lemma then implies that $Q_N = P_N$. Thus there is only one monic polynomial of degree N that is orthogonal to all polynomials of lower degree. This implies that the matrix H_N^* is nonsingular. \square

Theorem 10. *The matrix H_{N+k}^* is singular for $k = 1, 2, \dots$*

Proof: Instead of the basis of the monomials $\{z^p\}_{p \geq 0}$ we consider the basis $\{\psi_p(z)\}_{p \geq 0}$ where $\psi_p(z) := z^p$ for $p = 0, 1, \dots, N-1$ and $\psi_{N+p}(z) := z^p P_N(z)$ for $p = 0, 1, 2, \dots$. Let

$$F_l^* := \left[\langle \psi_p, \psi_q \rangle_\star \right]_{p,q=0}^{l-1}$$

be the corresponding $l \times l$ Gram matrix for $l = 1, 2, \dots$. Equation (16) then implies that $\det F_{N+k}^* = 0$ for $k = 1, 2, \dots$. One can easily verify that $\det F_l^* = \det H_l^*$ for $l = 1, 2, \dots$. This proves the theorem. \square

We have now identified $P_N(z)$ as the regular FOP of degree N and we have shown that regular FOPs of degree larger than N do not exist. Note that s_0^* is equal to the sum of the residues of $1/f$ at the points Z_1, \dots, Z_N and hence it is not necessarily different from zero. Therefore, the regular FOP of degree 1 with respect to the form $\langle \cdot, \cdot \rangle_\star$ does not always exist, in contrast to $\langle \cdot, \cdot \rangle$.

The zero/eigenvalue properties discussed in the previous section hold not only for $\langle \cdot, \cdot \rangle$ but for every symmetric bilinear form. The zeros Z_1, \dots, Z_N can therefore also be calculated by solving a generalized eigenvalue problem. The following result can be proved in the same way as Theorem 3. Let $H_k^{* <}$ be the Hankel matrix

$$H_k^{* <} := \begin{bmatrix} s_1^* & s_2^* & \cdots & s_k^* \\ s_2^* & & \ddots & \vdots \\ \vdots & \ddots & & \vdots \\ s_k^* & \cdots & \cdots & s_{2k-1}^* \end{bmatrix}$$

for $k = 1, 2, \dots$.

Theorem 11. *The eigenvalues of the pencil $H_N^{* <} - \lambda H_N^*$ are given by Z_1, \dots, Z_N .*

In the next section we will compare the accuracy obtained via Theorem 3 to the accuracy obtained via Theorem 11.

Algorithm ZEROS/FDF can be seen as a general algorithm for computing zeros of FOPs applied to the specific form $\langle \cdot, \cdot \rangle$. The following theorem leads to an analogous algorithm for the form $\langle \cdot, \cdot \rangle_\star$ and hence for the zeros Z_1, \dots, Z_N . We will call this approach ‘‘algorithm ZEROS/F.’’

Let $\{\varphi_t^*\}_{t \geq 0}$ denote the FOPs associated with $\langle \cdot, \cdot \rangle_\star$. Define the matrices G_k^* and $G_k^{\star z}$ as

$$G_k^* := \left[\langle \varphi_r^*, \varphi_s^* \rangle \right]_{r,s=0}^{k-1} \quad \text{and} \quad G_k^{\star z} := \left[\langle \varphi_r^*, z\varphi_s^* \rangle \right]_{r,s=0}^{k-1}$$

for $k = 1, 2, \dots$. The following results can be proved in the same way as Theorem 4 and Corollary 5.

Theorem 12. *Let $t \geq 1$ be a regular index for $\langle \cdot, \cdot \rangle_\star$ and let $z_{t,1}^*, \dots, z_{t,t}^*$ be the zeros of the regular FOP φ_t^* . Then the eigenvalues of the pencil $G_t^{\star z} - \lambda G_t^*$ are given by $z_{t,1}^*, \dots, z_{t,t}^*$.*

Corollary 13. *The eigenvalues of $G_N^{\star z} - \lambda G_N^*$ are given by Z_1, \dots, Z_N .*

If instead of N only an upper bound for N is available, then the value of N can be computed via a stopping criterion similar to the one discussed after Theorem 6.

5. Numerical Examples

In the following examples we have considered the case that γ is a circle. The computations have been done via Matlab 5 (with floating point relative accuracy $\approx 2.2204 \cdot 10^{-16}$).

The following integration algorithm is used to approximate the form $\langle \cdot, \cdot \rangle_\star$. Let γ be the circle with centre c and radius ρ . Then

$$\langle \phi, \psi \rangle_\star = \rho \int_0^1 \phi(c + \rho e^{2\pi i \theta}) \psi(c + \rho e^{2\pi i \theta}) \frac{1}{f(c + \rho e^{2\pi i \theta})} e^{2\pi i \theta} d\theta. \quad (17)$$

Since this is the integral of a periodic function over a complete period, the trapezoidal rule is an appropriate quadrature rule. If $F^* : [0, 1] \rightarrow \mathbb{C}$ is the integrand in the right-hand side of (17), then the q -point trapezoidal rule approximation to $\langle \phi, \psi \rangle_\star$ is given by

$$\langle \phi, \psi \rangle_\star = \int_0^1 F^*(\theta) d\theta \approx \frac{1}{q} \sum_{k=0}^q \prime\prime F^*(k/q) =: T_q.$$

The double prime indicates that the first and the last term of the sum are to be multiplied by $1/2$. As F^* is periodic with period one, we may rewrite T_q as

$$T_q = \frac{1}{q} \sum_{k=0}^{q-1} F^*(k/q).$$

This shows that T_q indeed depends on q (and not $q + 1$) points. As

$$T_{2q} = \frac{1}{2}T_q + T_{q \rightarrow 2q}$$

where

$$T_{q \rightarrow 2q} := \frac{1}{2q} \sum_{k=0}^{q-1} F^*\left(\frac{2k+1}{2q}\right),$$

successive doubling of q enables us in each step to reuse the integrand values needed in the previous step. In the following examples we started with $q = 16$ and continued doubling q until $|T_{2q} - T_q|$ was sufficiently small.

Table 1. Ordinary moments s_p and s_p^* in Example 1

p	s_p	s_p^*
0	6.0	$-7.5 \cdot 10^{-2}$
1	2.0	$2.8 \cdot 10^{-1}$
2	$-1.4 \cdot 10^1$	$-9.1 \cdot 10^{-1}$
3	$-8.5 \cdot 10^1$	2.1
4	$-3.0 \cdot 10^1$	-2.0
5	$7.0 \cdot 10^2$	2.4
6	$2.5 \cdot 10^3$	$-2.3 \cdot 10^1$
7	$-1.0 \cdot 10^3$	1.5 101
8	$-3.1 \cdot 10^4$	$9.9 \cdot 10^1$
9	$-7.6 \cdot 10^4$	$4.6 \cdot 10^2$
10	$1.3 \cdot 10^5$	$-4.8 \cdot 10^2$
11	$1.2 \cdot 10^6$	$-5.3 \cdot 10^3$

Example 1. Let $f(z) = e^{3z} - 2z \cos z - 1$. This function was also considered in [24]. All its zeros are simple. Suppose that γ is the circle $\gamma = \{z \in \mathbb{C} : |z| = 4\}$. Then $N = 6$. Let us try an approach based on ordinary moments. Table 1 contains approximations for $s_p = \langle 1, z^p \rangle$ and $s_p^* = \langle 1, z^p \rangle_*$ for $p = 0, 1, \dots, 11$. Note that in both cases the order of magnitude changes as p increases. The computed approximations for the zeros Z_1, \dots, Z_N obtained via Theorem 3 and 11 are shown in Table 2 and 3, respectively. The digits that are not correct are underlined. Observe that the approximations for the zeros are very accurate. Using ordinary moments has the advantage that only $2N$ integrals have to be calculated and hence, compared to algorithms ZEROS/FDF and ZEROS/F, the arithmetic cost is rather limited. Also, a significant part of the computation required for each integrand is the same for all of the integrands (namely, the computation of f'/f or $1/f$). By programming the quadrature algorithm in such a way that it is able to integrate a *vector* of similar integrals, these common calculations need be done

Table 2. Approximations for the zeros obtained via the ordinary moments s_p

$-2.186079491175828 \cdot 10^{-13}$	–	$i: 1.727623083122153 \cdot 10^{-12}$
$5.308949302929420 \cdot 10^{-1}$	+	$i: 1.331791876750615$
$5.308949302928376 \cdot 10^{-1}$	–	$i: 1.331791876751221$
-1.844233953262199	–	$i: 4.204494152042317 \cdot 10^{-14}$
1.414607177658190	+	$i: 3.047722062627169$
1.414607177658185	–	$i: 3.047722062627173$

Table 3. Approximations for the zeros obtained via the ordinary moments s_p^*

$5.879198486593449 \cdot 10^{-14}$	+	$i: 1.896836726398249 \cdot 10^{-14}$
$5.308949302929366 \cdot 10^{-1}$	+	$i: 1.331791876751066$
$5.308949302929205 \cdot 10^{-1}$	–	$i: 1.331791876751080$
-1.844233953262213	–	$i: 1.231196347425826 \cdot 10^{-15}$
1.414607177658181	+	$i: 3.047722062627166$
1.414607177658180	–	$i: 3.047722062627167$

Table 4. The number of correct significant digits in case $f(z) = J_0(z)$

Exact zeros	s_p^*	ZEROS/F	s_p	ZEROS/FDF
2.404825557695773	5	12	6	12
5.520078110286311	2	11	4	10
8.653727912911013	2	10	4	9
11.79153443901428	3	11	5	9
14.93091770848778	3	11	5	9
18.07106396791092	3	11	4	10
21.21163662987926	4	11	5	11
24.35247153074930	5	11	6	11
27.49347913204025	7	11	7	12

only once for each integrand evaluation point. However, as the following example shows, ordinary moments do not always lead to such accurate results.

Example 2. The Wilkinson polynomial and also functions that have clusters of zeros are typical, although somewhat extreme, examples where an approach based on ordinary moments is likely to fail. The following function is another example. Suppose that $f(z) = J_0(z)$, the Bessel function of the first kind and of order zero. It is known that this function has only positive real zeros and that all these zeros are simple (see, e.g., Watson [29]). In that sense it is related to the Wilkinson polynomial. Suppose that $\gamma = \{z \in \mathbb{C} : |z - 15| = 14.5\}$. Then $N = 9$.

Table 4 gives for each zero the number of correct significant digits obtained via the ordinary moments s_p^* (Theorem 11), algorithm ZEROS/F, the ordinary moments s_p (Theorem 3) and algorithm ZEROS/FDF.

Observe that the approximations for the zeros obtained via algorithms ZEROS/F and ZEROS/FDF are more accurate than the approximations obtained via ordinary moments. Of course, there is clearly a trade-off between obtained accuracy and cost. We advise the reader to start with the cheapest approach, i.e., the approach based on the ordinary moments s_p^* . If the computed approximations for the zeros are not sufficiently accurate to be refined via an iterative method (one that does not need the derivative, of course), then one can use algorithm ZEROS/F or switch to one of the approaches that use both f and f' .

Example 3. Let us illustrate how the stopping criterion of algorithm ZEROS/F can be used to determine the value of N in case only an upper bound for N is known. Consider again the function $f(z) = e^{3z} - 2z \cos z - 1$ and suppose that $\gamma = \{z \in \mathbb{C} : |z| = 5\}$. Then $N = 7$. Let us assume that only the upper bound 20 is known. Algorithm ZEROS/F defines the FOP φ_1^* as an inner polynomial and φ_2^* as a regular FOP. At this point the algorithm asks itself whether N is equal to two. It computes $|\langle \varphi_2^*, \varphi_2^* \rangle_\star|$. To take into account the accuracy lost during the evaluation of the quadrature formula, this quantity is scaled in a certain way. (We omit the details of this heuristic strategy. See [24].) The resulting floating point number is given by

$$1.998545018990362,$$

which is certainly not “sufficiently small” (we use 10^{-8} as a threshold) and hence the algorithm continues. It defines φ_3^* as an inner polynomial and φ_4^* as a regular FOP. Then it checks if N is equal to four. It compares

$$1.981687581683116$$

to 10^{-8} and continues. The polynomial φ_5^* is defined as a regular FOP. The algorithm again decides to continue and defines φ_6^* as a regular FOP. The corresponding floating point number is given by

$$0.3794164188056766$$

and the algorithm continues. It defines φ_7^* as a regular FOP. We have now reached the actual value of N . The scaled counterparts of the inner products that correspond to the sequence (14) are given by

$$\begin{aligned} &9.190814944765118 \cdot 10^{-16} \\ &1.799485800789563 \cdot 10^{-15} \\ &4.008700446099430 \cdot 10^{-15} \\ &5.548436809880727 \cdot 10^{-15} \end{aligned}$$

$$\begin{aligned}
& 6.603511781861113 \cdot 10^{-15} \\
& 5.538342691314587 \cdot 10^{-15} \\
& 3.494634208761963 \cdot 10^{-15} \\
& 4.116380174988637 \cdot 10^{-16} \\
& 5.379567405602837 \cdot 10^{-15} \\
& 8.806423950940129 \cdot 10^{-15} \\
& 8.912210606016112 \cdot 10^{-15} \\
& 5.866528582137192 \cdot 10^{-15} \\
& 1.127215921207880 \cdot 10^{-15}
\end{aligned}$$

and hence the algorithm decides that N is equal to seven and it stops. The computed approximations for the zeros are given by

$$\begin{array}{rcl}
-2.212860324230451 \cdot 10^{-11} & + & i : 5.610894531592185 \cdot 10^{-12} \\
\underline{5.308949303037738} \cdot 10^{-1} & + & i : \underline{1.331791876751059} \\
\underline{5.308949303027991} \cdot 10^{-1} & - & i : \underline{1.331791876755293} \\
-\underline{1.844233953258748} & - & i : 1.244550552500599 \cdot 10^{-12} \\
\underline{1.414607177657119} & + & i : \underline{3.047722062626751} \\
\underline{1.414607177657241} & - & i : \underline{3.047722062626826} \\
-\underline{4.603562881675490} & + & i : 3.443757237606488 \cdot 10^{-14}
\end{array}$$

The correct significant digits are underlined. Let us now compare this with the approach based on ordinary moments. The following theorem generalizes Theorem 11.

Theorem 14. *Let t be an integer $\geq N$. The eigenvalues of the pencil $H_t^{*\leq} - \lambda H_t^*$ are given by the zeros Z_1, \dots, Z_N and $t - N$ eigenvalues that may assume arbitrary values.*

Proof: Instead of the basis of the monomials $\{z^p\}_{p \geq 0}$ we consider again the basis $\{\psi_p(z)\}_{p \geq 0}$ where $\psi_p(z) := z^p$ for $p = 0, 1, \dots, N-1$ and $\psi_{N+p}(z) := z^p P_N(z)$ for $p = 0, 1, 2, \dots$, cf. the proof of Theorem 10. Define

$$F_t^{*\leq} := \left[\langle \psi_p, z \psi_q \rangle_\star \right]_{p,q=0}^{t-1} \quad \text{and} \quad F_t^* := \left[\langle \psi_p, \psi_q \rangle_\star \right]_{p,q=0}^{t-1}.$$

Then one can easily show that the generalized eigenvalue problem $H_t^{*\leq} x = \lambda H_t^* x$ is equivalent to the problem $F_t^{*\leq} y = \lambda F_t^* y$. Here $y := U_t^{-1} x$ where U_t denotes the

unit upper triangular matrix that contains the coefficients (in the standard monomial basis) of the polynomials $\psi_0(z), \psi_1(z), \dots, \psi_{t-1}(z)$. Equation (16) then implies that

$$F_t^{* <} = \begin{bmatrix} H_N^{* <} & 0 \\ 0 & 0 \end{bmatrix} \quad \text{and} \quad F_t^* = \begin{bmatrix} H_N^* & 0 \\ 0 & 0 \end{bmatrix}.$$

This proves the theorem. \square

Each of these indeterminate generalized eigenvalues corresponds to two corresponding zeros on the diagonals of the generalized Schur decomposition of the Hankel matrices $H_t^{* <}$ and H_t^* . When actually calculated, these diagonal entries are different from zero because of roundoff errors, and Matlab returns their quotient as an eigenvalue. Thus, by solving the 20×20 generalized eigenvalue problem $H_{20}^{* <} - \lambda H_{20}^*$ we obtain approximations for the seven zeros Z_1, \dots, Z_N and 13 spurious eigenvalues. The latter can be detected by evaluating f at the computed eigenvalues and also by taking into account that the computed approximations for the zeros are likely to lie inside γ or at least quite close to it. The approximations for the zeros obtained in this way are 1 to 3 digits less accurate than the approximations obtained via algorithm ZEROS/F. By solving the 7×7 generalized eigenvalue problem, one obtains approximations that are about as accurate as those computed by algorithm ZEROS/F.

Example 4. Let us consider a function that has multiple zeros. Suppose that

$$f(z) = z^2(z - 2)^2[e^{2z} \cos z + z^3 - 1 - \sin z]$$

and let $\gamma = \{z \in \mathbb{C} : |z| = 3\}$. Note that f has a triple zero at the origin and a double zero at $z = 2$. The total number of zeros of f that lie inside γ is equal to eight, $N = 8$. By using the ordinary moments s_p^* we obtain the following approximations for the zeros:

1.183531315599526 10^{-4}	–	i : 8.840648137844101 10^{-7}
–5.994094794302300 10^{-5}	–	i : 1.020522445441414 10^{-4}
–5.841218335364184 10^{-5}	+	i : 1.029363093460768 10^{-4}
<u>2.000000113260292</u>	+	i : 9.253727402009306 10^{-7}
<u>1.999999886743732</u>	–	i : 9.253736788382785 10^{-7}
– <u>4.607141197285995</u> 10^{-1}	+	i : <u>6.254277693471380</u> 10^{-1}
– <u>4.607141197287246</u> 10^{-1}	–	i : <u>6.254277693472881</u> 10^{-1}
<u>1.664682869740608</u>	+	i : 1.093307455221265 10^{-12}

The correct significant digits are underlined. Algorithm ZEROS/F gives comparable results. Note how the obtained accuracy diminishes as the multiplicity of the zero increases.

Acknowledgements

The authors would like to thank Bernard Mourrain for very stimulating discussions. This research was partially supported by projects #G.0278.97 “Orthogonal Systems and their Applications” and #G.0261.96 “Counting and Computing all Isolated Solutions of Systems of Nonlinear Equations” of the Fund for Scientific Research–Flanders (FWO–Vlaanderen).

References

- [1] Bojanczyk, A. W., Heinig, G.: A multi-step algorithm for Hankel matrices. *J. Complexity* *10*, 142–164 (1994).
- [2] Botten, L. C., Craig, M. S., McPhedran, R. C.: Complex zeros of analytic functions. *Comput. Phys. Commun.* *29*, 245–259 (1983).
- [3] Boulton, T., Sikorski, K.: An optimal complexity algorithm for computing the topological degree in two dimensions. *SIAM J. Sci. Stat. Comput.* *10*, 686–698 (1989).
- [4] Bultheel, A., Van Barel, M.: *Linear algebra, rational approximation and orthogonal polynomials*. Amsterdam: North-Holland 1997.
- [5] Cabay, S., Meleshko, R.: A weakly stable algorithm for Padé approximants and the inversion of Hankel matrices. *SIAM J. Matrix Anal. Appl.* *14*, 735–765 (1993).
- [6] Carpentier, M. C., Dos Santos, A. F.: Solution of equations involving analytic functions. *J. Comput. Phys.* *45*, 210–220 (1982).
- [7] Davies, B.: Locating the zeros of an analytic function. *J. Comput. Phys.* *66*, 36–49 (1986).
- [8] Delves, L. M., Lyness, J. N.: A numerical method for locating the zeros of an analytic function. *Math. Comput.* *21*, 543–560 (1967).
- [9] Draux, A.: *Polynômes orthogonaux formels—applications*. Berlin Heidelberg New York: Springer 1983.
- [10] Draux, A.: Formal orthogonal polynomials revisited. *Applications. Numer. Alg.* *11*, 143–158 (1996).
- [11] Freund, R. W., Zha, H.: A look-ahead algorithm for the solution of general Hankel systems. *Numer. Math.* *64*, 295–321 (1993).
- [12] Gautschi, W.: On the construction of Gaussian quadrature rules from modified moments. *Math. Comput.* *24*, 245–260 (1970).
- [13] Gautschi, W.: On generating orthogonal polynomials. *SIAM J. Sci. Stat. Comput.* *3*, 289–317 (1982).
- [14] Gautschi, W.: On the sensitivity of orthogonal polynomials to perturbations in the moments. *Numer. Math.* *48*, 369–382 (1986).
- [15] Gautschi, W.: How (un)stable are Vandermonde systems? In: *Asymptotic and computational analysis: conference in honor of Frank W. J. Olver’s 65th Birthday* (Wong, R., ed.), pp. 193–210. New York: Marcel Dekker 1990.
- [16] Gautschi, W., Inglese, G.: Lower bounds for the condition number of Vandermonde matrices. *Numer. Math.* *52*, 241–250 (1988).
- [17] Gragg, W. B., Gutknecht, M. H.: Stable look-ahead versions of the Euclidean and Chebyshev algorithms. In: *Approximation and computation: A festschrift in honor of walter Gautschi* (Žahar, R. V. M., ed.), pp. 231–260. Basel: Birkhäuser 1994.
- [18] Gutknecht, M. H.: A completed theory of the unsymmetric Lanczos process and related algorithms, Part I. *SIAM J. Matrix Anal. Appl.* *13*, 594–639 (1992).
- [19] Gutknecht, M. H.: A completed theory of the unsymmetric Lanczos process and related algorithms, Part II. *SIAM J. Matrix Anal. Appl.* *15*, 15–58 (1994).
- [20] Henrici, P.: *Applied and computational complex analysis: I. Power series—integration—conformal mapping—location of zeros*. New York: Wiley 1974.
- [21] Herlocker, J., Ely, J.: An automatic and guaranteed determination of the number of roots of an analytic function interior to a simple closed curve in the complex plane. *Reliable Comput.* *1*, 239–250 (1995).
- [22] Ioakimidis, N. I.: Quadrature methods for the determination of zeros of transcendental functions—a review. In: *Numerical integration: recent developments, software and applications* (Keast, P., Fairweather, G., eds.), pp. 61–82. Dordrecht: Reidel 1987.
- [23] Ioakimidis, N. I., Anastasselou, E. G.: A modification of the Delves-Lyness method for locating the zeros of analytic functions. *J. Comput. Phys.* *59*, 490–492 1985.

- [24] Kravanja, P., Sakurai, T., Van Barel, M.: On locating clusters of zeros of analytic functions. *BIT* 39, (1999) (to appear).
- [25] Kravanja, P., Van Barel, M., Ragos, O., Vrahatis, M. N., Zafiroopoulos, F. A.: ZEAL: A mathematical software package for computing zeros of analytic functions. *Comput. Phys. Commun.* (to appear).
- [26] Li, T.-Y.: On locating all zeros of an analytic function within a bounded domain by a revised Delves/Lyness method. *SIAM J. Numer. Anal.* 20, 865–871 (1983).
- [27] Lloyd, N. G.: Degree theory. Cambridge: Cambridge University Press 1978.
- [28] Traub, J. F.: Wasilkowski, G. W., Woźniakowski, H.: Information-Based Complexity. New York: Academic Press, 1988.
- [29] Watson, G. N.: A treatise on the theory of Bessel functions. Cambridge: Cambridge University Press 1966.
- [30] Wilkinson, J. H.: The evaluation of the zeros of ill-conditioned polynomials. Part I. *Numer. Math.* 1, 150–166 (1959).
- [31] Ying, X.: A reliable root solver for automatic computation with application to stress analysis of a composite plane wedge. D. Sc. dissertation, Washington University in St. Louis, 1986.
- [32] Ying, X., Katz, I. N.: A reliable argument principle algorithm to find the number of zeros of an analytic function in a bounded domain. *Numer. Math.* 53, 143–163 (1988).

P. Kravanja
Katholieke Universiteit Leuven
Department of Computer Science
Celestijnenlaan 200 A
B-3001 Heverlee
Belgium
e-mail: Peter.Kravanja@na-net.ornl.gov

M. Van Barel
Katholieke Universiteit Leuven
Department of Computer Science
Celestijnenlaan 200 A
B-3001 Heverlee
Belgium
e-mail: Marc.VanBarel@cs.kuleuven.ac.be