2013

TR-2013006: Polynomial Evaluation and Interpolation and Transformations of Matrix Structures

Victor Y. Pan

Polynomial Evaluation and Interpolation and Transformations of Matrix Structures

Victor Y. Pan

Department of Mathematics and Computer Science
Lehman College and the Graduate Center of the City University of New York
Bronx, NY 10468 USA
victor.pan@lehman.cuny.edu,
home page: http://comet.lehman.cuny.edu/vpan/

Abstract. Multipoint polynomial evaluation and interpolation are fundamental for modern numerical and symbolic computing. The known algorithms solve both problems over any field of constants in nearly linear arithmetic time, but the cost grows to quadratic for numerical solution. We decrease this cost dramatically and for a large class of inputs yield nearly linear time as well. We first restate our tasks as multiplication of a Vandermonde matrix and its inverse by a vector, then transform this matrix into other structured matrices, and finally apply a variant of the Multipole celebrated techniques to achieve the desired speedup for the computations with polynomials, Vandermonde matrices and their transposes. An important impact of our work is a new demonstration of the power of the method of the transformation of matrix structures, which we proposed in [P90]. At the end we comment on further applications and extension of this method to computations with structured matrices, polynomials, and rational functions.

Key words: Polynomials, Rational functions, Multipoint evaluation, Interpolation, Vandermonde matrices, Cauchy matrices, Transformation of matrix structures, Multipole method

AMS Subject Classification: 12Y05, 15A04, 47A65, 65D05, 68Q25

1 Introduction

Multipoint polynomial evaluation and interpolation are fundamental for modern numerical and symbolic computations. Both problems can be solved in nearly linear arithmetic time over any field [F72], [H72], [MB72], [BP94], [P01], [GG03]. For numerical solution, in the presence of rounding errors, however, these algorithms are prone to error propagation, and in spite of some research advances, say, in [PSLT93], [P95], and [PZHY97], the users employ quadratic time algorithms (cf. [KZ08], [BF00], [P64], [BP70], [BEGO08]). We propose new numerical algorithms that accelerate the known algorithms dramatically and for a large class of inputs solve both problems also in nearly linear arithmetic time.

We first employ the known equivalence of our tasks to multiplication of a Vandermonde matrix and its inverse by a vector and then propose a novel FFT-based reduction of these computations to operations with HSS matrices. ("HSS" is the acronym for "hierarchically semiseparable".) Finally we apply the efficient numerical algorithms of [MRT05], [CGS07], and [XXG12], based on the Multipole celebrated techniques of [GR87] and [CGR98]. Overall this enables dramatic acceleration of the known numerical algorithms, and for a large class of inputs we yield nearly linear arithmetic time for approximate multipoint polynomial evaluation and interpolation as well as for multiplication by vectors of a Vandermonde matrix and its transpose and inverse (see Theorem 15).

Our advance can be extended to computations with some popular classes of matrices having structures of Vandermonde and Cauchy types and is a new demonstration of the power of transformation of matrix structures, proposed in [P90]. Previous combinations of this method with the Multipole techniques enabled efficient approximate numerical solution of a nonsingular Toeplitz linear system of equations by using nearly linear arithmetic time [MRT05], [CGS07], and [XXG12]. We extend the power of this ingenious work to a broader class of structured matrices, including Vandermonde matrices, and in Section 9 and [Pa] we comment on further applications and extension of this method to computations with polynomials, structured matrices and rational functions.

We organize our paper as follows. In the next three sections we recall some definitions and basic results for general, banded, block diagonal and block tridiagonal matrices as well as the evaluation and interpolation tasks for polynomials and rational functions and fast FFT-based reduction of these tasks to each other and to multiplication of Vandermonde and Cauchy matrices and their inverses by a vector. In Section 5 we cover some HSS techniques. In Section 6 we approximate Cauchy matrices by HSS matrices, and in Section 7 we devise our nearly optimal algorithms. In Section 8 we extend these algorithms to a wider class of matrices based on their displacement representation. We refer the reader to [Pa] on a more comprehensive treatment of this subject. We conclude the paper with Section 9, where we summarize our study and suggest some natural research directions.

2 Definitions and auxiliary results

Hereafter "flop" stands for "arithmetic operation performed in the field $\mathbb{C}$ of complex numbers with no error". $|\mathcal{S}|$ denotes the cardinality of a set $\mathcal{S}$.

$M = (m_{i,j})_{i,j=0}^{m-1,n-1}$ is an $m \times n$ matrix. M^T is its transpose, M^H is its Hermitian transpose. For two sets $\mathcal{I} \subseteq \{1,\dots,m\}$ and $\mathcal{J} \subseteq \{1,\dots,n\}$, define the submatrices $M(\mathcal{I},\mathcal{J}) = (m_{i,j})_{i\in\mathcal{I},j\in\mathcal{J}}$. Write $M(\mathcal{I},.) = M(\mathcal{I},\mathcal{J})$ where $\mathcal{J} = \{1,\dots,n\}$. Write $M(.,\mathcal{J}) = M(\mathcal{I},\mathcal{J})$ where $\mathcal{I} = \{1,\dots,m\}$. $\mathcal{C}(B)$ and $\mathcal{R}(B)$ are the sets of indices of the row and column sets of a submatrix B of $M = (m_{i,j})_{i,j=1}^{m,n}$, respectively. $\mathcal{R}(B) = \mathcal{I}$ and $\mathcal{C}(B) = \mathcal{J}$ if and only if $B = M(\mathcal{I},\mathcal{J})$. An $m \times n$ matrix M has a non-unique *generating pair* (F, G^T) *of a length* ρ if

$M = FG^T$ for two matrices $F \in \mathbb{C}^{m\times\rho}$ and $G \in \mathbb{C}^{n\times\rho}$. The rank of a matrix is the minimum length of its generating pairs.

Theorem 1. *A matrix M has a rank ρ if and only if it has a nonsingular $\rho\times\rho$ submatrix $M(\mathcal{I},\mathcal{J})$, and if so, then $M = M(.,\mathcal{J})M(\mathcal{I},\mathcal{J})^{-1}M(\mathcal{I},.)$, that is $(M(.,\mathcal{J}), M(\mathcal{I},\mathcal{J})^{-1}M(\mathcal{I},.)$ and $(M(.,\mathcal{J}), M(\mathcal{I},\mathcal{J})^{-1}M(\mathcal{I},.)$ are generating pairs of M of length ρ.*

We refer to generating pairs and *triples* such as $(M(.,\mathcal{J}), M(\mathcal{I},\mathcal{J})^{-1}, M(\mathcal{I},.))$ as *generators*.

$(B_0 \mid \ldots \mid B_{k-1})$ is a $1\times k$ block matrix with k blocks $B_0,\ldots,B_{k-1}$, whereas $D_B = \mathrm{diag}(B_0,\ldots,B_{k-1}) = \mathrm{diag}(B_j)_{j=0}^{k-1}$ is a $k\times k$ block diagonal matrix with k diagonal blocks $B_0,\ldots,B_{k-1}$, possibly rectangular. For 1×1 blocks $b_j = B_j$ we arrive at a vector $\mathbf{b}^T = (b_0 \mid \ldots \mid b_{k-1})$ and a $k\times k$ diagonal matrix $D_{\mathbf{b}} = \mathrm{diag}(b_j)_{j=0}^{k-1}$, respectively. $O = O_{m,n}$ is the $m\times n$ matrix filled with zeros. $I = I_n = \mathrm{diag}(1)_{j=0}^{n-1}$ is the $n\times n$ identity matrix. An $n\times n$ matrix M is nonsingular if $\mathrm{rank}(M) = n$ or equivalently if it has the inverse $X = M^{-1}$ such that $XM = MX = I$. M is a $k\times l$ *unitary* matrix if $M^HM = I_l$ or $MM^H = I_k$. These two equations imply one another and imply that $M^H = M^{-1}$ if $k = l$.

$\alpha(M)$ and $\beta(M)$ denote the numbers of flops required for computing the vectors $M\mathbf{u}$ and $M^{-1}\mathbf{u}$, respectively, maximized over all vectors $\mathbf{u}$ and minimized over all algorithms, and we write $\beta(M) = \infty$ where the matrix M is singular. The straightforward algorithms support the following bound.

Theorem 2. *$\alpha(M) \le (2m + 2n - 1)\rho - m < 2(m + n)\rho$ for an $m\times n$ matrix M given with its generating pair of a length ρ.*

$||M|| = ||M||_2$ denotes the 2-norm of a matrix M. It holds that $||U|| = ||V|| = 1$ and $||MV|| = ||UM|| = ||M||$ for any unitary matrices U and V. For an $m\times n$ matrix M of a rank ρ we define its *SVD* or *full SVD*, $M = S_M\Sigma_MT_M^H$ where S_M and T_M are square orthogonal matrices, $S_MS_M^H = S_M^TS_M = I_m$, $T_MT_M^H = T_M^HT_M = I_n$, $\Sigma_M = \mathrm{diag}(\mathrm{diag}(\sigma_j(M))_{j=1}^{\rho}, O_{m-\rho,n-\rho})$, $\sigma_j = \sigma_j(M) = \sigma_j(M^H)$ is the jth largest *singular value* of a matrix M for $j = 1,\ldots,\rho$, $\sigma_j = 0$ for $j > \rho$, $\sigma_\rho > 0$, $\sigma_1 = \max_{||\mathbf{x}||=1}||M\mathbf{x}|| = ||M||$, and

$$\min_{\mathrm{rank}(B)\le s-1} ||A - B|| = \sigma_s(A),\ \ s = 1,2,\ldots \tag{1}$$

Theorem 3. *Assume a 3×3 block matrix $M = \begin{pmatrix} U & O & X\\ V & B & Y\\ W & O & Z\end{pmatrix}$ with ϵ-rank at most ρ and with $\sigma_\rho(B) > \epsilon$. Then the matrix $M_- = \begin{pmatrix} U & X\\ W & Z\end{pmatrix}$ has norm at most ϵ.*

Proof. Equation (1) for $s = \rho + 1$ implies that $\mathrm{rank}(M_\epsilon) \le \rho$ for some ϵ-perturbation of the matrix M. It follows that this perturbation annihilates the matrix M_- because $\sigma_\rho(B) > \epsilon$, and so the matrix Σ_B has ϵ-rank at least ρ.

Delete the $n-\rho$ last columns of the matrices S_M and Σ_M and the $m-\rho$ last rows of the matrices Σ_M and T_M^H and obtain *compact SVD* $M = \bar{S}_M \bar{\Sigma}_M \bar{T}_M^H$, defining a generating triple $(\bar{S}_M, \bar{\Sigma}_M, \bar{T}^H)$ of the minimum length ρ for the matrix M. See [S98, Section 5.1] on the perturbation study of SVDs. Generating triples of the minimum length for a given matrix can be also supplied by its less costly rank revealing factorizations such as ULV and URV factorizations in [CGS07], [XXG12], and [XXCBa], where the matrices U and V are unitary, whereas the matrices L and R are triangular.

A matrix $\tilde{M}$ is an ϵ-*approximation* of a matrix M if $||\tilde{M} - M|| \le \epsilon$. We similarly define other ϵ-concepts such as the ϵ-*rank* of a matrix M, denoting the integer $\min_{||\tilde{M}-M||\le\epsilon} \operatorname{rank}(\tilde{M})$ (with the implication that $\sigma_\rho(M) > \epsilon \ge \sigma_{\rho+1}(M)$), an ϵ-*basis* for a linear space $\mathbb{S}$ of dimension k, denoting a set of vectors that ϵ-approximate the k vectors of a basis for this space, and an ϵ-generator of a matrix, which is a generator of its ϵ-approximation. $\alpha_\epsilon(M)$ and $\beta_\epsilon(M)$ replace the bounds $\alpha(M)$ and $\beta(M)$ where we ϵ-approximate the vectors $M\mathbf{u}$ and $M^{-1}\mathbf{u}$ instead of evaluating them. The *numerical rank* is the ϵ-rank for small ϵ. A matrix is *ill conditioned* if its rank exceeds its numerical rank.

Theorem 4. *(See [S98, Corollary 1.4.19] for $P = -M^{-1}E$.) Suppose M and $M+E$ are two nonsingular matrices of the same size and $||M^{-1}E|| = \theta < 1$. Then $||I - (M+E)^{-1}M|| \le \frac{\theta}{1-\theta}$ and $||(M+E)^{-1} - M^{-1}|| \le \frac{\theta}{1-\theta}||M^{-1}||$. In particular $||(M+E)^{-1} - M^{-1}|| \le 0.5\ ||M^{-1}||$ if $\theta \le 1/3$.*

3 Banded, block banded, and block diagonal matrices

A matrix $B = (b_{ij})_{i,j=1}^{m,n}$ has a lower bandwidth l, an upper bandwidth u, and less than $(l+u+1)\min\{m,n\}$ nonzero entries if $b_{ij} = 0$ unless $-l \le j-i \le u$. Such matrix is called *banded* if $l+u$ is small in context.

Theorem 5. *(Cf. [GL96].) It holds that $\alpha(B) = O((l+u+1)(m+n))$ and if the matrix B is nonsingular, then $\beta(B) = O((l+u)^2 n)$.*

An $n \times n$ lower triangular banded matrix B with a lower bandwidth l such that $n = kl$ can be expressed as a block bidiagonal matrix $B = (B_{p,q})_{p,q=0}^{k-1}$ with $l \times l$ blocks $B_{p,q}$ where $B_{p,q} = O$ unless $0 \le p-q \le 1$. (To extend the block bidiagonal representation to the case where $n = (k-1)l + g$ and $0 < g < l$, embed the matrix into a proper $(kl) \times (kl)$ block bidiagonal matrix or allow a $g \times g$ block $B_{0,0}$ or $B_{k-1,k-1}$.) If the matrix B is nonsingular, then the block diagonal matrix $D_B = \operatorname{diag}(B_q)_{q=0}^{k-1}$ is nonsingular as well, all diagonal blocks of the block bidiagonal matrix $B' = D_B^{-1}B$ are equal to the identity matrix I_l, and so the matrix $S = I_n - B'$ is filled with zeros, except for its first block subdiagonal. One can obtain such matrix by moving all block rows of a block diagonal matrix $-\operatorname{diag}(S_q)_{q=0}^{k-1}$ one block down. Note that $S^k = O$, obtain that $\widehat{B}^{-1} = (I_n - S)^{-1} = I + S + \cdots + S^{k-1}$, and explicitly express the inverse, as we show in the 5×5 case next. (See some generalizations in Section 5.)

$$\widehat{B}^{-1} = \begin{pmatrix} I & O & O & O & O \\ S_0 & I & O & O & O \\ S_1 S_0 & S_1 & I & O & O \\ S_2 S_1 S_0 & S_2 S_1 & S_2 & I & O \\ S_3 S_2 S_1 S_0 & S_3 S_2 S_1 & S_3 S_2 & S_3 & I \end{pmatrix}. \tag{2}$$

If $m \times n$ matrices M and $D = \mathrm{diag}(D_q)_{q=0}^{k-1}$ share the diagonal blocks $D_0, \dots, D_{k-1}$, then we write $D = \mathrm{diag}(M)$, call the ordered set $\widehat{D} = (D_0, \dots, D_{k-1})$ the *block diagonal* of the matrix M, and call the matrix $N(M) = M - \mathrm{diag}(M)$ its *neutered complement* (cf. [MRT05, Section 1]). In Section 6 we deal with block tridiagonal matrices extended into the southeastern and northwestern corners, respectively, as we show in the following 5×5 example (cf. Figure 1),

$$T = \begin{pmatrix} \Sigma_0 & B_0 & O & O & A_0 \\ A_1 & \Sigma_1 & B_1 & O & O \\ O & A_2 & \Sigma_2 & B_2 & O \\ O & O & A_3 & \Sigma_3 & B_3 \\ B_4 & O & O & A_4 & \Sigma_4 \end{pmatrix}. \tag{3}$$

We also represent such *extended tridiagonal* matrices as *extended block diagonal* matrices in two dual ways, $T = \Sigma^{(c)} = \mathrm{diag}(\Sigma_0^{(c)}, \dots, \Sigma_{k-1}^{(c)}) = \Sigma^{(r)} = \mathrm{diag}(\Sigma_0^{(r)}, \dots, \Sigma_{k-1}^{(r)})$. Here $\Sigma_q^{(c)} = \begin{pmatrix} B_{q-1 \mod k} \\ \Sigma_q \\ A_{q+1 \mod k} \end{pmatrix}$ and $\Sigma_q^{(r)} = (A_q \mid \Sigma_q \mid B_q)$ for $q = 0, \dots, k-1$ denote the *extended diagonal blocks*, each made up of a triple of the blocks of the matrix T that form adjacent chains if we glue together the lower and upper boundaries of the matrix as well as its right and left boundaries.

4 Dense structured matrices. Polynomial and rational evaluation and interpolation. DFT, IDFT, FFT, IFFT, and some transformations of matrix structures

Write $T = (t_{i-j})_{i,j=0}^{n-1}$, $H = (h_{i+j})_{i,j=0}^{n-1}$, $V = V_{\mathbf{s}} = (s_i^j)_{i,j=0}^{n-1}$, and $C = C_{\mathbf{s},\mathbf{t}} = \left(\frac{1}{s_i - t_j}\right)_{i,j=0}^{n-1}$ to denote *Toeplitz, Hankel, Vandermonde, and Cauchy* matrices, respectively, which are four classes of most popular dense structured matrices, each having n^2 entries defined by at most $2n$ parameters. Every matrix of these classes as well as its inverse (if defined) can be multiplied by a vector in nearly linear arithmetic time, and proper displacement operators enable extension of these properties to more general classes of matrices having structures of Toeplitz, Hankel, Vandermonde and Cauchy types (see Section 7). Here are some sample links of computations with structured matrices to fundamental polynomial and rational computations (see more on that in [P01, Chapters 2 and 3]).

Problem 1. Multipoint polynomial evaluation or Vandermonde-by-vector multiplication.

INPUT: $2n$ complex scalars $p_0, \dots, p_{n-1}; s_0, \dots, s_{n-1}$.
OUTPUT: n complex scalars $v_0, \dots, v_{n-1}$ satisfying

$$v_i = p(s_i) \text{ for } p(x) = p_0 + p_1 x + \dots + p_{n-1}x^{n-1} \text{ and } i = 0, \dots, n-1 \tag{4}$$

or equivalently

$$V\mathbf{p} = \mathbf{v} \text{ for } V = V_{\mathbf{s}} = (s_i^j)_{i,j=0}^{n-1},\ \mathbf{p} = (p_j)_{j=0}^{n-1},\ \text{and } \mathbf{v} = (v_i)_{i=0}^{n-1}. \tag{5}$$

Problem 2. Polynomial interpolation or the solution of a Vandermonde linear system of equations.
INPUT: $2n$ complex scalars $v_0, \dots, v_{n-1}; s_0, \dots, s_{n-1}$, the last n of them distinct.
OUTPUT: n complex scalars $p_0, \dots, p_{n-1}$ satisfying equations (4) and (5).

Problem 3. Multipoint rational evaluation or Cauchy-by-vector multiplication.
INPUT: $3n$ complex scalars $s_0, \dots, s_{n-1}; t_0, \dots, t_{n-1}; v_0, \dots, v_{n-1}$.
OUTPUT: n complex scalars $v_0, \dots, v_{n-1}$ satisfying

$$v_i = \sum_{j=0}^{n-1} \frac{u_j}{s_i - t_j} \text{ for } i = 0, \dots, n-1 \tag{6}$$

or equivalently

$$C\mathbf{u} = \mathbf{v} \text{ for } C = C_{\mathbf{s},\mathbf{t}} = \left(\frac{1}{s_i - t_j}\right)_{i,j=0}^{n-1},\ \mathbf{u} = (u_j)_{j=0}^{n-1},\ \text{and } \mathbf{v} = (v_i)_{i=0}^{n-1}. \tag{7}$$

Problem 4. Rational interpolation or the solution of a Cauchy linear system of equations.
INPUT: $3n$ complex scalars $s_0, \dots, s_{n-1}; t_0, \dots, t_{n-1}; v_0, \dots, v_{n-1}$, the first $2n$ of them distinct.
OUTPUT: n complex scalars $u_0, \dots, u_{n-1}$ satisfying (6) and (7).

Every rational function $v(x) = \frac{p(x)}{t(x)}$, for $p(x)$ of equation (4), $t(x) = \prod_{j=0}^{n-1}(x - t_j)$ and n distinct knots $t_0, \dots, t_{n-1}$, can be represented as $v(x) = \sum_{j=0}^{n-1} \frac{u_j}{x - t_j}$, which turns into equations (6) if we write $v_i = v(s_i)$ for $i = 0, \dots, n-1$.

If the knots $s_0, \dots, s_{n-1}, t_0, \dots, t_{n-1}$ are distinct, then the matrices $V = V_{\mathbf{s}}$ and $C = C_{\mathbf{s},\mathbf{t}}$ are nonsingular, and Problems 2 and 4 have unique solution because (see, e.g., [P01, Section 3.6])

$$\det V_{\mathbf{s}} = \prod_{i>j}(s_i - s_j),\ \ \det C_{\mathbf{s},\mathbf{t}} = \prod_{i<j}(s_j - s_i)(t_i - t_j) / \prod_{i,j}(s_i - t_j). \tag{8}$$

How many flops do we need for solving Problems 1–4? Horner's algorithm of 1819 evaluates the polynomial $p(x)$ of (4) at a single knot $x = s$ by using $2n - 2$ flops, and this is optimal [P66]. (The algorithm was used by Newton in 1669, by a number of medieval mathematicians, and in the Chinese Nine Chapters on the Mathematical Art (the Han Dynasty, 202 BC – 220 AD).) For Problem 1 with n

knots, n applications of Horner's algorithm involve $2(n-1)n$ flops, but this is not optimal anymore. The algorithms of [F72], [GGS87], and [MB72] solve Problems 1–3 by using $O(n\log^2(n)\log(\log(n)))$ flops over any field of constants, which is within a factor of $\log(n)\log(\log(n))$ from the optimum [S73], [B-O83]. Equation (9) of the next section extends this cost bound to Problem 4. For numerical solution of Problems 1–4, however, the users employ quadratic time algorithms to avoid error propagation (cf. [KZ08], [BF00], [P64], [BP70], [BEGO08]).

Numerical solution in nearly linear arithmetic time $O(n\log(n))$ is known, however, for the important special case of these problems where the knots are the roots of 1. Write $\omega = \omega_n = \exp(2\pi\sqrt{-1}/l)$ (to denote a primitive nth root of 1), $s_i = \omega^i$ for $i = 0,\dots,n-1$, $V_{\mathbf{s}} = (\omega^{ij})_{i,j=0}^{n-1}$, and $\Omega = \frac{1}{\sqrt{n}}V_{\mathbf{s}}$. Then $\Omega^H\Omega = I_n$, $\Omega^T = \Omega$, Ω and $\Omega^H = \Omega^{-1} = \frac{1}{\sqrt{n}}(\omega^{-ij})_{i,j=0}^{n-1}$ are unitary matrices, and Problems 1 and 2 turn into the computational problems of the forward and inverse discrete Fourier transforms (hereafter *DFT* and *IDFT*). The *FFT* (Fast Fourier transform) and Inverse FFT (*IFFT*) are numerically stable algorithms that perform DFT and IDFT by using $1.5n\log_2(n)$ and $1.5n\log_2(n)+n$ flops, respectively, if n is a power of 2 (cf. [BP94, Sections 1.2 and 3.4]), whereas Generalized FFT and IFFT use $O(n\log(n))$ flops to perform DFT and IDFT for any n [P01, Problem 2.4.2].

In spite of some common properties the four matrix structures have quite distinct features. The matrix structure of Cauchy type is invariant in row and column interchange (in contrast to the structures of Toeplitz and Hankel types) and enables expansion of the matrix entries into Loran's series (unlike the structures of the three other types). The paper [P90], however, links the four structures to each other by means of structured matrix multiplication and exploits *this link to extend all successful matrix inversion algorithms for the matrices of any of the four classes to the matrices of the three other classes.* The present paper shows a new specialization of this general approach.

Our progress in numerical solution of Problems 1 and 2 relies on linking together Vandermonde and Cauchy matrix computations. In particular recall from [P01, Section 3.6] that

$$C_{\mathbf{s},\mathbf{t}} = \operatorname{diag}(t(s_i)^{-1})_{i=0}^{n-1} V_{\mathbf{s}} V_{\mathbf{t}}^{-1} \operatorname{diag}(t'(t_j))_{j=0}^{n-1} \tag{9}$$

where $\mathbf{s} = (s_i)_{i=0}^{n-1}$, $\mathbf{t} = (t_i)_{i=0}^{n-1}$, and $t(x) = \prod_{i=0}^{n-1}(x-t_i)$.

Remark 1. One can compute the values $-t_0^n \dots, -t_{n-1}^n$ by using $O(n\log n)$ flops, and then the computation of the coefficients of the polynomial $v(x) = t(x) - x^n$ given its values $v(t_i) = -t_i^n$ for $i = 0,\dots,n-1$ turns into a special case of Problem 2 of polynomial interpolation.

Next note that $C_{\mathbf{s},\mathbf{t}}^T = -C_{\mathbf{t},\mathbf{s}}$, $\Omega^T = \Omega$, and $D^T = D$ for a diagonal matrix D and deduce from the above equation that

$$V_{\mathbf{s}} = \operatorname{diag}(t(s_i))_{i=0}^{n-1} C_{\mathbf{s},\mathbf{t}} \operatorname{diag}(t'(t_j)^{-1})_{j=0}^{n-1} V_{\mathbf{t}}, \tag{10}$$

$$V_{\mathbf{s}}^T = -V_{\mathbf{t}}^T \operatorname{diag}(t'(t_j)^{-1})_{j=0}^{n-1} C_{\mathbf{t},\mathbf{s}} \operatorname{diag}(t(s_i))_{i=0}^{n-1}, \tag{11}$$

$$V_{\mathbf{s}}^{-1} = V_{\mathbf{t}}^{-1} \operatorname{diag}(t'(t_j))_{j=0}^{n-1} C_{\mathbf{s},\mathbf{t}}^{-1} \operatorname{diag}(t(s_i)^{-1})_{i=0}^{n-1}, \tag{12}$$

$$V_{\mathbf{s}}^{-T} = -\operatorname{diag}(t(s_i)^{-1})_{i=0}^{n-1} C_{\mathbf{t},\mathbf{s}}^{-1} \operatorname{diag}(t'(t_j))_{j=0}^{n-1} V_{\mathbf{t}}^{-T}. \tag{13}$$

For $\mathbf{t} = (f\omega^j)_{j=0}^{n-1}$, the knots t_i are the roots of 1 scaled by f, $t(x) = x^n - f^n$, $t'(x) = nx^{n-1}$, $V_{\mathbf{t}} = \sqrt{n}\Omega \operatorname{diag}(f^j)_{j=0}^{n-1}$, and $V_{\mathbf{t}}^{-1} = \operatorname{diag}(f^{-j})_{j=0}^{n-1} \Omega^H/\sqrt{n}$, and then equations (9)–(13) take the following form (where we use that $\Omega^T = \Omega$),

$$C_{\mathbf{s},f} = \sqrt{n} \operatorname{diag}\left(\frac{1}{s_i^n - f^n}\right)_{i=0}^{n-1} V_{\mathbf{s}} \operatorname{diag}(f^{-j})_{j=0}^{n-1} \Omega^H \operatorname{diag}(\omega^{-j})_{j=0}^{n-1}, \tag{14}$$

$$V_{\mathbf{s}} = \frac{1}{\sqrt{n}} \operatorname{diag}\left(s_i^n - f^n\right)_{i=0}^{n-1} C_{\mathbf{s},f} \operatorname{diag}(\omega^j)_{j=0}^{n-1} \Omega \operatorname{diag}(f^j)_{j=0}^{n-1}, \tag{15}$$

$$V_{\mathbf{s}}^T = -\frac{1}{\sqrt{n}} \operatorname{diag}(f^j)_{j=0}^{n-1} \Omega \operatorname{diag}(\omega^j)_{j=0}^{n-1} C_{f,\mathbf{s}} \operatorname{diag}(s_i^n - f^n)_{i=0}^{n-1}, \tag{16}$$

$$V_{\mathbf{s}}^{-1} = \sqrt{n} \operatorname{diag}(f^{-j})_{j=0}^{n-1} \Omega^H \operatorname{diag}(\omega^{-j})_{j=0}^{n-1} C_{\mathbf{s},f}^{-1} \operatorname{diag}\left(\frac{1}{s_i^n - f^n}\right)_{i=0}^{n-1}, \tag{17}$$

$$V_{\mathbf{s}}^{-T} = -\sqrt{n} \operatorname{diag}\left(\frac{1}{s_i^n - f^n}\right)_{i=0}^{n-1} C_{f,\mathbf{s}}^{-1} \operatorname{diag}(\omega^{-j})_{j=0}^{n-1} \Omega^H \operatorname{diag}(f^{-j})_{j=0}^{n-1}. \tag{18}$$

The latter equations link Vandermonde matrices and their transposes, inverses, and transposes of the inverses to the Cauchy matrices with the knot set $\mathcal{T} = \{t_j = f\omega^j,\ j = 0, \dots, n-1\}$, which we call *CV matrices* and denote $C_{\mathbf{s},f}$. Their transposes have the knot set $\mathcal{S} = \{s_j = f\omega^j,\ j = 0, \dots, n-1\}$, are linked to transposed Vandermonde matrices, said to be *CV^T matrices*, and are denoted $C_{f,\mathbf{t}}$. [P01, Equation (3.4.1)] links together a Vandermonde matrix, its transpose, inverse and the inverse of the transpose. Here are some other sample links among matrix structures, more comprehensively covered in [P01, Sections 4.7 and 4.8] and [Pa].

Theorem 6. *(i) JH and HJ are Toeplitz matrices if H is a Hankel matrix, and vice versa. (ii) $H = V^TV = (\sum_{k=0}^{m-1} s_k^{i+j})_{i,j=0}^{n-1}$ is a Hankel matrix for any $m \times n$ Vandermonde matrix $V = (s_i^j)_{i,j=0}^{m-1,n-1}$.*

5 HSS matrices and neutered blocks

We are going to link CV matrices to structured matrices that extend banded matrices and their inverses. We refer the reader to [EGHa], [EGHb], [VVGM05], [MRT05], [CGS07], [VVM07], [VVM08], [X12], and the bibliography therein on this extension and its variations under the names of matrices with low Hankel rank, rank structured matrices, quasiseparable, and weakly, recursively, or sequentially semiseparable matrices. We cite [GR87], [CGR98], [LRT79], [P93], and [PR93] on the related subjects of Multipole and Nested Dissection Algorithms, and their parallel implementation.

Definition 1. *(See Figures 2 and 3.) Assume an* $m \times n$ *matrix* M *with a block diagonal* $\widehat{\Sigma}' = \{\Sigma'_q\}_{q=0}^{k-1}$ *and extended block diagonals* $\widehat{\Sigma}^{(c)} = (\Sigma_0^{(c)}, \dots, \Sigma_{k-1}^{(c)})$ *and* $\widehat{\Sigma}^{(r)} = (\Sigma_0^{(r)}, \dots, \Sigma_{k-1}^{(r)})$. *Here we use the notation of the end of Section 3; hereafter for simplicity we also write* $\widehat{\Sigma} = (\Sigma_0, \dots, \Sigma_{k-1})$ *instead of* $\widehat{\Sigma}^{(c)} = (\Sigma_0^{(c)}, \dots, \Sigma_{k-1}^{(c)})$, *dropping the superscript* $^{(c)}$. *(i) (See [MRT05, Section 1].) A block of the matrix* M *is* neutered *unless it overlaps the extended block diagonal, and so such block is either* subdiagonal *or* superdiagonal. *For a set of consecutive indices* $\mathcal{J}$ *remove all rows of the block column* $M(., \mathcal{J})$ *that overlap the extended block diagonal and obtain the* neutered block column $N(\mathcal{J}) = N^{(c)}(\mathcal{J})$. *(ii) It is a* basic neutered block column, dual *to an extended diagonal block* Σ_q *and denoted* N_q, *if* $\mathcal{J} = \mathcal{C}(\Sigma_q)$, *that is if this neutered block column shares its column indices with the submatrix* Σ_q. *(iii) The* neutered union $N(N(\mathcal{J}), N(\mathcal{K}))$ *of two neutered block columns* $N(\mathcal{J})$ *and* $N(\mathcal{K})$ *is the neutered block column* $N(\mathcal{J} \cup \mathcal{K})$. *(iv) Similarly define neutered block rows* $N^{(r)}(\mathcal{I})$, *their unions, and the basic neutered block rows* $N_p^{(r)}$ *dual to the extended diagonal blocks* $\Sigma_p^{(r)}$.

Definition 2. *The matrix* M, *given with its extended block diagonal is an* (l, u)*-HSS block matrix if* l *is the maximum rank of its subdiagonal neutered blocks and if* u *is the maximum rank of its superdiagonal neutered blocks. The matrix* M *is* ρ-neutered *if it is a* (ρ, ρ)*-HSS block matrix or equivalently if* ρ *is the maximum rank of its neutered blocks. By replacing ranks with* ϵ*-ranks we extend this definition and introduce* (ϵ, l, u)*-HSS block matrices and* (ϵ, ρ)*-neutered matrices.*

We immediately verify the following results.

Theorem 7. *Every neutered block is a block submatrix of the neutered union of some basic neutered block columns as well as of some basic neutered block rows.*

Theorem 8. *Assume a matrix* $M = (B_0 \mid \dots \mid B_{k-1})$ *having* k *block columns* $B_0, \dots, B_{k-1}$ *and given with its extended block diagonal* $\widehat{\Sigma} = (\Sigma_0, \dots, \Sigma_{k-1})$ *and with* k *generating pairs of lengths at most* ρ *defining the* k *basic neutered block columns* $N_0, \dots, N_{k-1}$. *Then one can modify the extended diagonal blocks* Σ_q *to obtain* k *generating pairs of lengths at most* ρ *f defining the* k *block columns* B'_q *of the resulting matrix* $M' = (B'_0 \mid \dots \mid B'_{k-1})$.

Corollary 1. *Under the assumptions of Theorem 8 let* $\Sigma = \mathrm{diag}(M)$ *be generic matrix. Then* $\alpha(M) \leq \alpha(\Sigma) + (2m + 2n - 1)k\rho$.

Proof. Theorem 8 defines a generating pair of a length at most $k\rho$ for the matrix M', and so $\alpha(M') \leq (2m + 2n - 1)k\rho - m$ by virtue of Theorem 2. Furthermore the matrices M and M' share the entries of all basic neutered block columns, and so $\alpha(M - M') \leq \alpha(\Sigma)$. Combine the above estimates with the simple bound $\alpha(M) \leq \alpha(M') + \alpha(M - M') + m$.

The following theorem, extending Theorem 5, yields superior estimates in the case of ρ-neutered matrices and (l, u)-HSS block matrices M.

Theorem 9. *(See [EG02, Section 3], [CGS07, Sections 3 and 4], and our Remark 2.) Assume an (l,u)-HSS block matrix M of size $m \times n$ having an extended block diagonal $\widehat{\Sigma} = (\Sigma_0, \dots, \Sigma_{k-1})$ with $m_q \times n_q$ extended diagonal blocks Σ_q, $q = 0, \dots, k-1$. Then it holds that*
(i) $\alpha(M) \le 2\sum_{q=0}^{k-1}((m_q+n_q)(l+u)+m_qn_q+l^2k+u^2k) = O((l+u)(m+n))$, and (ii) if $m = n$ and if the matrix M is nonsingular, then $\beta(M) = O((l+u)^2n)$.

Part (i) is supported by the algorithms of [EG02, Section 3]. They separately multiply by a vector the extended block diagonal matrix $\Sigma = \mathrm{diag}(M)$, the subdiagonal (lower triangular) and superdiagonal (upper triangular) parts of the matrix M. Part (ii) is supported by the algorithms of [CGS07]. Both papers as well as and the study in [VVM07], [VVM08], [XXG12], [EGHa] and [EGHb] rely on the representation of an (l,u)-HSS matrix M with *HSS generators*, which we next demonstrate by a 4×4 example (cf. (2)) and then define in a theorem,

$$M = \begin{pmatrix} \Sigma_0 & S_0T_1 & S_0B_1T_2 & S_0B_1B_2T_3 \\ P_1Q_0 & \Sigma_1 & S_1T_2 & S_1B_2T_3 \\ P_2A_1Q_0 & P_2Q_1 & \Sigma_2 & S_2T_3 \\ P_3A_2A_1Q_0 & P_3A_2Q_1 & P_3Q_2 & \Sigma_2 \end{pmatrix}.$$

Theorem 10. *(Cf. [EGHa], [VVM07], [X12], the bibliography therein, and our Table 1 and Remark 2.) Assume a $k \times k$ matrix M with an extended block diagonal $\widehat{\Sigma} = (\Sigma_0, \dots, \Sigma_{k-1})$, where $\Sigma_q = M(I_q, J_q)$, $q = 0, \dots, k-1$. Then M is an (l,u)-HSS block matrix if and only if there exists a nonunique family of* HSS generators *$\{P_i,\ Q_h,\ S_h,\ T_i,\ A_g,\ B_g\}$ such that $M(\mathcal{I}_i, \mathcal{J}_h) = P_iA_{i-1} \cdots A_{h+1}Q_h$ and $M(\mathcal{I}_h, \mathcal{J}_i) = S_hB_{h+1} \cdots B_{i-1}T_i$ for $0 \le h < i < k$. Here P_i, Q_h and A_g are $|\mathcal{I}_i| \times l_i$, $l_{h+1} \times |\mathcal{J}_h|$, and $l_{g+1} \times l_g$ matrices, respectively, S_h, T_i and B_g are $|\mathcal{I}_h| \times u_{h+1}$, $u_i \times |J_i|$, and $u_g \times u_{g+1}$ matrices, respectively, $g = 1, \dots, k-2$, $h = 0, \dots, k-2$, and $i = 1, \dots, k-1$, $l = \max_g\{l_g\}$ and $u = \max_h\{u_h\}$.*

Table 1. The sizes of the HSS generators in Theorem 10

P_i	Q_h	A_g	S_h	T_i	B_g
$\lvert\mathcal{I}_i\rvert \times l_i$	$l_{h+1} \times \lvert\mathcal{J}_h\rvert$	$l_{g+1} \times l_g$	$\lvert\mathcal{I}_h\rvert \times u_{h+1}$	$u_i \times \lvert J_i\rvert$	$u_g \times u_{g+1}$

Based on this theorem one can redefine the (l,u)-HSS block matrices as the ones allowing representation with some HSS generator families $\{P_h,\ Q_i,\ A_g\}$ and $\{S_h,\ T_i,\ B_g\}$ whose *lower and upper lengths* or *orders* are equal to l and u, respectively.

Remark 2. The cited bibliography covers matrices with block diagonals, but the study presented there, including Theorems 9 and 10, is readily extended to the case of matrices with extended block diagonals. Moreover by replacing ranks with

ϵ-ranks, generators with ϵ-generators, computation with ϵ-approximation, $\alpha(M)$ with $\alpha_\epsilon(M)$, and $\beta(M)$ with $\beta_\epsilon(M)$, we extend the study, including Theorems 9 and 10, to the case of (ϵ, l, u)-HSS block matrices and (ϵ, ρ)-neutered matrices.

6 ϵ-approximation of CV matrices by HSS block matrices

[MRT05, Section 4] and [CGS07, Section 2.4] ϵ-approximate the matrix $C_{1,\omega_{2n}}$ by an (l, u)-HSS block matrix for $l = u = O(\log(n))$, ϵ of order $c'/n^{c''}$, and two positive constants c' and c'', and then solve a linear system of equations with this matrix numerically in nearly linear time by applying Multipole techniques. We yield similar results for multiplication by a vector of CV and CV^T matrices as well as of their inverses where they exist.

6.1 Small-rank approximation of Cauchy matrices where the knot sets $\mathcal{S}$ and $\mathcal{T}$ are separated from one another

In this subsection we closely follow [CGS07].

Definition 3. *(See [CGS07, page 1254].) For a* separation bound $\theta < 1$ *and a complex* separation center c*, two complex points s and t are* (θ, c)-separated *from one another if* $|\frac{t-c}{s-c}| \le \theta$*. Two sets of complex numbers $\mathcal{S}$ and $\mathcal{T}$ are* (θ, c)-separated *from one another if every two points $s \in \mathcal{S}$ and $t \in \mathcal{T}$ are (θ, c)-separated from one another. $\delta_{c,\mathcal{S}} = \min_{s\in\ \mathcal{S}} |s - c|$ and $\delta_{c,\mathcal{T}} = \min_{t\in\mathcal{T}} |t - c|$ denote the distances from the center c to the sets $\mathcal{S}$ and $\mathcal{T}$, respectively.*

Lemma 1. *(See [R85] and [CGS07, equation (2.8)].) Suppose two complex values s and t are (θ, c)-separated from one another for a positive $\theta < 1$ and a complex c and write $q = \frac{t-c}{s-c}$, $|q| \le \theta$. Then for every positive integer ρ it holds that*

$$\frac{1}{s-t} = \frac{1}{s-c}\sum_{i=0}^{\rho-1}\frac{(t-c)^i}{(s-c)^i} + \frac{q_\rho}{s-c} \text{ where } |q_\rho| = \frac{|q|^\rho}{1-|q|} \le \frac{\theta^\rho}{1-\theta}. \tag{19}$$

Proof. $\frac{1}{s-t} = \frac{1}{s-c}\frac{1}{1-q}$, $\frac{1}{1-q} = \sum_{i=0}^{\infty} q^i = (\sum_{i=0}^{\rho-1} q^i + \sum_{i=\rho}^{\infty} q^i) = (\sum_{i=0}^{\rho-1} q^i + \frac{q^\rho}{1-q})$.

Corollary 2. *(Cf. [CGS07, Section 2.2].) Suppose two sets of $2n$ distinct complex numbers $\mathcal{S} = \{s_0, \dots, s_{n-1}\}$ and $\mathcal{T} = \{t_0, \dots, t_{n-1}\}$ are (θ, c)-separated from one another for $0 < \theta < 1$ and a complex c. Define the Cauchy matrix $C = (\frac{1}{s_i - t_j})_{i,j=0}^{n-1}$ and write $\delta = \delta_{c,\mathcal{S}} = \min_{i=0}^{n-1} |s_i - c|$. Then for every positive integer ρ it is sufficient to use $2\rho n$ flops to compute the $n \times \rho$ matrices F and G such that*

$$F = (1/(s_i - c)^{\nu+1})_{i,\nu=0}^{n-1,\rho-1}, \ G = ((t_j - c)^\nu)_{j,\nu=0}^{n-1,\rho-1}, \tag{20}$$

$$C = FG^T + E, \quad ||E|| \le \frac{n\theta^\rho}{(1-\theta)\delta}. \tag{21}$$

Proof. Apply (19) for $s = s_i$, $t = t_j$, and all pairs (i, j) to deduce (20) and (21).

In the corollary and throughout we can replace $\delta = \delta_{c,\mathcal{S}} = \min_{i=0}^{n-1} |s_i - c|$ by $\delta = \delta_{c,\mathcal{T}} = \min_{j=0}^{n-1} |t_j - c|$ because of the symmetric roles of the sets $\mathcal{S}$ and $\mathcal{T}$.

Remark 3. Corollary 2 defines a $||E||$-generating pair of a length at most ρ for the Cauchy matrix C. Unless any of the values $1 - \theta$ and δ is small, the norm $||E||$ of (21) is small already for moderately large integers ρ. This implies small upper bounds on the numerical rank of a Cauchy matrix $C = (\frac{1}{s_i - t_j})_{i,j=0}^{n-1}$ whose parameter sets $\mathcal{S} = \{s_0, \ldots, s_{n-1}\}$ and $\mathcal{T} = \{t_0, \ldots, t_{n-1}\}$ are (θ, c)-separated from one another for some center c provided neither of the values $1 - \theta$ and δ is small. If such separation holds just for two subsets of the sets $\mathcal{S}$ and $\mathcal{T}$ that define a $k \times l$ Cauchy submatrix, then the numerical rank of the matrix C cannot exceed $2n - k - l + \rho$, and this implies that the matrix C is ill conditioned if it is nonsingular (cf. (8)) and if $k + l > n + \rho$. Now recall that equations (10)–(13) link Vandermonde matrices to CV and CV^T matrices via FFT and conclude that nonsingular Vandermonde matrices $V_{\mathbf{s}}$ are ill conditioned, except for their narrow subclass where the knots of the sets $\mathbb{S} = \{s_0, \ldots, s_{n-1}\}$ closely approximate all or almost all knots of the set $\{f\omega_n^i\}_{i=0}^{n-1}$ of the scaled nth roots of 1 for $|f| = 1$ (cf. [GI88]). For $n > 2$ this narrow subclass includes no Vandermonde matrices $V_{\mathbf{s}}$ with the sets $\mathcal{S}$ of real knots, because all real values are well separated from the roots of 1 outside quite small neighborhoods of 1 and -1.

6.2 An extended block diagonal of a CV matrix

Given a vector $\bar{\mathbf{s}} = (\bar{s}_i)_{i=0}^{n-1}$ and a complex point f on the unit circle $\{f : |f| = 1\}$, we seek a permutation matrix P and a ρ-neutered matrix that approximates the CV matrix $C = C_{\mathbf{s},f} = (\frac{1}{s_i - f\omega^{j-1}})_{i,j=0}^{n-1}$ within a norm bound ϵ for $\mathbf{s} = (s_i)_{i=0}^{n-1} = P\bar{\mathbf{s}}$ and $\rho = O(\log(n/\sqrt{\epsilon}))$, proceeding as follows.

Begin with defining a permutation matrix P and an extended block diagonal of the matrix C (cf. (3) and Figures 1 and 4). Assume the polar coordinates for the knots $\bar{s}_i = |\bar{s}_i| \exp(2\pi\bar{\phi}_i\sqrt{-1})$ where $0 \le \bar{\phi}_i < 2\pi$, $\bar{\phi}_i = 0$ if $\bar{s}_i = 0$, and $i = 0, \ldots, n-1$. Write $\phi_0 = \min_{i=0}^{n-1} \bar{\phi}_i$, reorder the angles $\bar{\phi}_i$ in the nondecreasing order breaking ties arbitrarily, let P denote the permutation matrix that defines this reordering, and then write $(\phi_i)_{i=0}^{n-1} = P(\bar{\phi}_i)_{i=0}^{n-1}$ and $(s_i)_{i=0}^{n-1} = P(\bar{s}_i)_{i=0}^{n-1}$. To simplify the notation assume that $f = 1$ and $n = hk$ for two positive integers h and k. Let $\mathcal{S}'_q$ and $\mathcal{T}_q = \{\omega^j\}_{j=qh}^{(q+1)h-1}$ denote the subsets of the sets $\mathcal{S}$ and $\mathcal{T}$, respectively, consisting of the knots s_i and $t_j = \omega^j$ such that

$$qh \le \phi_i < (q+1)h - 1, \; qh \le j < (q+1)h - 1, \; q = 0, \ldots, k-1. \tag{22}$$

The two subsets are the intersections of the sets $\mathcal{S}$ and $\mathcal{T}$ with the semi-open sector Γ_q of the complex plane bounded by the pairs of rays from the origin to the points ω^{hq} and $\omega^{(q+1)h}$. Write

$$\mathcal{S}_q = \mathcal{S}'_{q-1 \mod k} \cup \mathcal{S}'_q \cup \mathcal{S}'_{q+1 \mod k}, \; q = 0, \ldots, k-1. \tag{23}$$

Now for the matrix $C = \left(\frac{1}{s_i-\omega^j}\right)_{i,j=0}^{n-1}$ define the k diagonal blocks

$$\Sigma'_q = \left(\frac{1}{s_i - \omega^j}\right)_{i\in\mathcal{S}'_q,\ j\in\mathcal{T}_q},\ q = 0,\dots,k-1,$$

which have hn entries overall, and the k extended diagonal blocks

$$\Sigma_q = \left(\frac{1}{s_i - \omega^j}\right)_{i\in\mathcal{S}_q, j\in\mathcal{T}_q},\ q = 0,\dots,k-1, \tag{24}$$

which have $3hn$ entries overall, and then for every q, $q = 0,\dots,k-1$, partition the block column $\left(\frac{1}{s_i-\omega^j}\right)_{0\le i<n, j\in\mathcal{T}_q}$ into the block Σ_q and the dual basic neutered block column $N_q = N_q^{(c)}$. The following algorithm summarizes these computations.

Algorithm 1. Extended block diagonal.

INPUT: *three integers h, $k > 1$, and $n = hk$, and n complex values $\bar{s}_0,\dots,\bar{s}_{n-1}$ represented in polar coordinates.*

OUTPUT: *an $n\times n$ permutation matrix P, the vector $\mathbf{s} = (s_i)_{i=0}^{n-1} = P(\bar{s}_i)_{i=0}^{n-1}$, the CV matrix $C_{\mathbf{s},1} = (\frac{1}{s_i-\omega^j})_{i,j=0}^{n-1}$, the sets $\mathcal{S}_q \subseteq \mathcal{S} = \{s_0,\dots,s_{n-1}\}$ for $q = 0,\dots,k-1$, and an extended block diagonal matrix $\Sigma = \operatorname{diag}(\Sigma_q)_{q=0}^{k-1}$ whose extended diagonal blocks Σ_q satisfy equations (22)–(24) for the sets $\mathcal{T}_q = \{\omega^j\}_{j=qh}^{(q+1)h-1}$, $q = 0,\dots,k-1$.*

COMPUTATIONS:

1. *Reorder the knots $\bar{s}_i = |\bar{s}_i|\exp(2\pi\bar{\phi}_i\sqrt{-1})$ for $i = 0,\dots,n-1$ in nondecreasing order of the angles $\bar{\phi}_i$. Break ties arbitrarily. Output the permutation matrix P of this reordering and the vector $\mathbf{s} = (s_i)_{i=0}^{n-1} = P(\bar{s}_i)_{i=0}^{n-1}$.*
2. *Based on equations (22) compute the sets $\mathcal{S}'_q \subseteq \mathcal{S} = \{s_i\}_{i=0}^{n-1}$ for $q = 0,\dots,k-1$.*
3. *Based on equations (23) combine the triples of the latter sets to define and to output the sets $\mathcal{S}_q$ for $q = 0,\dots,k-1$.*
4. *Output the blocks $\Sigma_0,\dots,\Sigma_{k-1}$ satisfying equations (24).*

6.3 The ϵ-ranks of basic neutered block columns

Theorem 11. *(Cf. Remarks 4–7 and Figure 4.) Assume a CV matrix C output by Algorithm 1. Then (i) the extended diagonal blocks $\Sigma_0,\dots,\Sigma_{k-1}$ together have exactly $3hn$ entries and (ii) for $|p-q \mod k| > 1$ the row index sets $\mathcal{R}(\Sigma_p)$ and $\mathcal{R}(\Sigma_q)$ have no overlap. Furthermore (iii) there are points $c_0,\dots,c_{k-1}$ on the unit circle $\{z: |z| = 1\}$ and at the distance of at least $0.5/(kn)$ from the set $\mathcal{S}$ such that the sets $\mathcal{S}_p$ and $\mathcal{T}_q$ are (θ, c_q)-separated from one another for $\theta = (1.5\pi/k)/\sin(3\pi/k)$ as long as $|p-q \mod k| > 1$.*

Proof. One can readily verify parts (i) and (ii). Let us prove part (iii). Let $\mathcal{A}(s,t)$ denote the arc of the unit circle $\{z: |z|=1\}$ with the end points s and t. For every q, $q=0,\dots,k-1$, choose a center c_q on the arc $\mathcal{A}(\omega_{4n}^{(4q+1)h}, \omega_{4n}^{(4q+3)h})$ at the distance at least $\frac{2}{kn}$ from the set $\mathcal{S}$ (as we require). This is possible because the set has exactly n elements. The arc has length π/k and shares the midpoint $\omega_{2n}^{(2q+1)h}$ with the arc $\mathcal{A}(\omega^{qh}, \omega^{(q+1)h})$ of length $2\pi/k$, on which all points of the set $\mathcal{T}_q$ lie. Therefore these points lie at the distance less than $\frac{3\pi}{2k}$ from the center c_q. Furthermore, unless $|p-q|\le 1$ or $|p-q|=k-1$, all points of the set $\mathcal{S}_p$ lie outside the sector of the complex plane bounded by the pairs of rays from the origin to the points $\omega^{(q-1)h}$ and $\omega^{(q+2)h}$, and then $\text{Distance}(c_q, \mathcal{S}_p) \ge \sin(3\pi/k)$. Therefore the sets $\mathcal{S}_p$ and $\mathcal{T}_q = \{\omega^j\}_{j=(q-1)h}^{hq-1}$ are (θ, c_q)-separated from one another for $\theta = (1.5\pi/k)/\sin(3\pi/k)$ unless $|p-q|\le 1$ or $|p-q|=k-1$. It holds that $1.5\pi/k \approx \sin(1.5\pi/k)$ for sufficiently large values $k=n/h$, whereas $\sin(1.5\pi/k)/\sin(3\pi/k) = 0.5/\cos(1.5\pi/k)$. By summarizing these estimates we obtain part (iii) of the theorem.

Corollary 3. *Suppose Algorithm 1 has output a CV matrix C and its extended block diagonal. Assume a positive integer ρ and write $\theta = \frac{1.5\pi/k}{\sin(3\pi/k)}$ and $\epsilon = \frac{n\theta^\rho}{(1-\theta)\sin(3\pi/k)}$. Then all basic neutered block columns have ϵ-ranks at most ρ.*

Proof. Theorem 11 implies that every basic neutered block column is (θ, c_q)-separated for some center c_q and for the above θ. Apply Corollary 2 for $\delta = \sin(3\pi/k)$ and deduce the claimed bound on the ϵ-ranks.

Remark 4. For a positive ϵ deduce from the corollary that $\frac{1}{\theta^\rho} = \frac{n}{\epsilon}\frac{1}{1-\theta}\ \frac{1}{\sin(3\pi/k)}$. Consequently, $\rho = (\log_2(\frac{n}{\epsilon}) + \log_2(\frac{1}{1-\theta}) + \log_2(\frac{1}{\sin(3\pi/k)})/\log_2(\frac{1}{\theta})$. Let us bound ρ from above in terms of k, n and ϵ. We can assume that n is reasonably large and then choose, say, $k\ge 12$ and obtain from Theorem 11 that $\theta < 0.5554$ (one can specify such bounds for other choices of k, noting that $\theta \to 1/2$ as $k = n/h \to \infty$). Now deduce that $\rho \le 1.2(\log_2(kn/\epsilon) - 2) < 1.2\log_2(kn/\epsilon)$ for $k \ge 12$. Finally substitute $k\le n$ and obtain

$$\rho \le \rho \le 1.2(\log_2(kn/\epsilon) - 2) < 2.4\ \log_2(n/\sqrt{\epsilon}),$$
$$\epsilon < kn/2^{\rho/1.2} \le n^2/2^{\rho/1.2} \text{ for } k \ge 2.$$

Remark 5. In the proof of Theorem 11 we can rotate all sectors Γ_q and centers c_q by a fixed angle ϕ, $0 \le \phi < 2\pi$, redefine the block diagonal matrix Σ accordingly, and then readily extend Theorem 11 and Corollary 3. In particular rotation by the angle π/k produces new basic neutered block columns, each overlapping a pair of old adjacent basic neutered block columns.

Remark 6. Among the basic neutered block columns N_q only N_0 and N_{k-1} are blocks, whereas for $q=1,\dots,k-2$ they are made up of the pairs of super- and subdiagonal blocks. By gluing the upper and lower boundaries of the matrix C, however, we can turn all these pairs into single blocks. Furthermore our study is invariant under block row and block column permutations, with which we can move any diagonal block into the northwestern position $(0,0)$ or the southeastern position $(k-1,k-1)$.

Remark 7. One can extend the results of this section in various ways. To simplify the notation we assumed that $f=1$ and $n=hk$ in Algorithm 1 and throughout, but the same arguments and proofs can be applied for any complex f and any triple of integers (h,k,n) satisfying $|f|=1$ and $k=\lceil n/h\rceil$. Since $-C_{\mathbf{s},\mathbf{t}}^T = C_{\mathbf{t},\mathbf{s}}$ one can replace the CV matrix of the corollary by a CV^T matrix $(\frac{1}{f\omega^i - t_j})_{i,j=0}^{n-1}$ for $|f|=1$. The proof techniques enable extension to rectangular CV and CV^T matrices as well as to a Cauchy matrix $C_{\mathbf{s},\mathbf{t}}$ with the set $\mathcal{S}$ more or less equally spaced about a segment of a line or a smooth curve on the complex plane, the unit circle $\{x: \ |x=1|\}$ being an example of such a curve. If the set $\mathcal{S}$ has such distribution on the complex plane, then the same argument bounds the ϵ-rank of all basic neutered block rows.

6.4 The $((2k-1)\epsilon)$-ranks of neutered blocks

Next, under some additional assumptions, we extend the bound ρ of Corollary 3 on the ϵ-rank from the basic neutered block columns N_j at first to the neutered unions of the pairs and the chains of such adjacent block columns and then (by virtue of Theorem 7) to all neutered blocks of the matrix C. To extend Corollary 3 to all neutered blocks of the CV matrices C we narrow the class of these matrices. The next definition narrows this class stronger than we need, but it helps introduce the subsequent definition that we really use.

Definition 4. *A matrix C is ϵ-*uniformly *ρ-neutered if its every $s\times s$ neutered block N has ϵ_+-rank $\rho \geq \min\{s,\rho\}$ for some $\epsilon_+ > \epsilon$.*

Definition 5. *(See Figures 5 and 6.) A 5-tuple (h,k,n,ρ,ϵ) is* valid *if h, k, n, and ρ are positive integers, $\epsilon > 0$, and the bounds of Remark 4 hold. For a CV matrix C of Corollary 3 define its two extended block diagonals, that is $\widehat{\Sigma}$, as in the corollary, and $\widehat{\Sigma}'$, as in Remark 5 for $\phi = \pi/k$. For $q=0,\dots,k-1$ let $N_q = N_c(\Sigma_q)$ and $N_q' = N_c(\Sigma_q')$ denote the basic neutered block columns dual to the diagonal blocks Σ_q and Σ_q' of these extended block diagonals, respectively, and let $\cap_q$ and $\cap_q'$ denote the two neutered blocks defined by the common entries of the matrix pairs (N_q, N_q') and $(N_{q+1 \mod k}, N_q')$, respectively. Then the matrix C is ϵ-*selectively *ρ-neutered (for this 5-tuple (h,k,n,ρ,ϵ)) if all blocks $\cap_q$ and $\cap_q'$ have ϵ_+-ranks at least ρ for some $\epsilon_+ > \epsilon$.*

Remark 8. For an *ϵ-selectively* ρ-neutered matrix we must have $\lfloor h/2\rfloor \geq \rho$ because each matrix $\cap_q$ and $\cap_q'$ has at least $\lfloor h/2\rfloor$ columns and has ϵ_+-rank at least ρ, but there is no such contradictions to the choice of $h=2\rho$ or $h=2\rho+1$ for a fixed integer ρ.

Theorem 12. *Assume an ϵ-*selectively *ρ-neutered $n\times n$ CV matrix C for a valid 5-tuple (h,k,n,ρ,ϵ). Then the neutered unions of all chains of adjacent basic neutered block columns N_q have $((2k-1)\epsilon)$-ranks at most ρ.*

Proof. By virtue of Corollary 3 the basic neutered block columns N_q and N'_q have ϵ-ranks at most ρ for all q, but actually even the blocks $\cap_q$ have ϵ-ranks exactly ρ for all q because C is assumed to be an ϵ-*selectively* ρ-neutered matrix. Let $\cap_q = \bar{S}_{\cap_q}\bar{\Sigma}_{\cap_q}\bar{T}^H_{\cap_q}$ be compact SVD for $0 \le q < k$. Then $\bar{S}^H_{\cap_q}\cap_q = \bar{\Sigma}_{\cap_q}\bar{T}^H_{\cap_q} = (M_\rho \mid E_q)^T$ where M_ρ is a $\rho \times \rho$ matrix, $||E_q|| \le \epsilon$, and where $\sigma_\rho(M_\rho) = \sigma_\rho(\cap_q) > \epsilon$ because the matrix is ϵ-*selectively* ρ-neutered. Represent the matrix $U_q = \bar{S}^H_{\cap_q}N(N_q, N'_q)$ as $\begin{pmatrix} W & M_\rho & Y \\ X & E_q & Z \end{pmatrix}$, apply Theorem 3 to each of the matrices $\begin{pmatrix} W & M_\rho \\ X & E_q \end{pmatrix}$ and $\begin{pmatrix} M_\rho & Y \\ E_q & Z \end{pmatrix}$, and deduce that $||(X \mid E_q)|| \le \epsilon$ and $||(E_q \mid Z)|| \le \epsilon$. Consequently only the first ρ rows of the matrix $U_q - E'_q$ are nonzero for some matrix E'_q satisfying $||E'_q|| \le 2\epsilon$. Extend this argument to prove that for any pair of integers l and q, $0 < l < k$, $0 \le q < k$, and the neutered union U of the chain

$$N(N_q, N'_q, N_{q+1 \mod k}, N'_{q+1 \mod k}, \ldots, N_{q+l-1 \mod k}, N_{q+l-1 \mod k})$$

of $2l-1$ overlapping neutered blocks, there is a matrix $E_{q,l}$ such that $||E_{q,l}|| \le (2l-1)\epsilon$ and the matrix $U - E_{q,l}$ has only ρ nonzero rows, and so $\text{rank}(U - E_{q,l}) \le \rho$. This implies the theorem.

Remark 9. The assumption of Theorem 12 that ϵ_+-ranks of all blocks $\cap_q$ and $\cap'_q$ are at least ρ for $q = 0, \ldots, k-1$ and any $\epsilon_+ > \epsilon$ can be quite readily verified for the much studied CV matrix $C_{1,\omega_{2n}}$. [MRT05] and [CGS07] state logarithmic bounds on its ϵ-rank but seem to omit detailed proofs.

Theorems 7 and 12 and Remark 2 together imply the following corollary.

Corollary 4. *A CV $n \times n$ matrix is $((2k-1)\epsilon, \rho)$-neutered if it is ϵ-selectively ρ-neutered for a valid 5-tuple $(h, k, n, \rho, \epsilon)$.*

7 Multiplication of *CV* and Vandermonde matrices and their transposes and inverses by a vector

Given an $n \times n$ CV matrix, we seek an $(n\epsilon)$-approximation of its product by a vector. Our algorithm supporting the following theorem accelerates by a factor of $\sqrt{n/\log(n)}$ the known algorithms provided that $\log(1/\epsilon) = O(\log(n))$.

Theorem 13. *Assume positive integers k and n, $n > k$, complex $s'_0, \ldots, s'_{n-1}$, and f, $|f| = 1$, positive ϵ, and a CV matrix $C' = (\frac{1}{s'_i - \omega^j})_{i,j=0}^{n-1}$. Then it holds that (i) $\alpha_{(k+1)\epsilon}(C') < 4n\sqrt{3n\rho}$ for ρ of Remark 4 and in particular $\alpha_{n\epsilon}(C') \le 4n\sqrt{7.2n\log_2(n/\sqrt{\epsilon})}$ for $\rho = 2.4\log_2(n/\sqrt{\epsilon})$ and $k < n$. (ii) Consequently it holds that $\alpha_{n\epsilon}(C') = O(n\sqrt{n\log(n)})$ where $\log(1/\epsilon) = O(\log(n))$.*

Proof. Apply Algorithm 1 to compute a permutation matrix P, an extended block diagonal matrix Σ, and the basic neutered block columns $N_0, \ldots, N_{k-1}$ of the matrix $C = PC'$. Clearly $\alpha_{(k+1)\epsilon}(C') = \alpha_{(k+1)\epsilon}(C) \leq \alpha_\epsilon(\Sigma) + \sum_{q=0}^{k-1} \alpha_\epsilon(N_q) + kn$. By virtue of Theorem 11, every basic neutered block column N_q has rank at most ρ and the matrix Σ has at most $3hn$ nonzero entries. Therefore, $\alpha(\Sigma) \leq 6hn - n$, whereas by virtue of Theorem 2, $\alpha(N_q) < 2(h + n - n_q)\rho - \rho$ for the $(n - n_q) \times h$ basic neutered block column N_q. Note that $\sum_{q=0}^{k-1} n_q = 3n$, and so $\sum_{q=0}^{k-1} \alpha(N_q) \leq 2(h+n)k\rho - 3n\rho - k\rho = 2kn\rho - (n+k)\rho$ because $hk = n$. Combine the above bounds and deduce that $\alpha_{(k+1)\epsilon}(C) \leq 6hn + 2kn\rho - (n + k)\rho - n < 6hn + 2kn\rho$. Choose $k\rho = 3h = 3n/k$, and deduce that $k = \sqrt{3n/\rho}$, $k\rho = \sqrt{3n\rho}$, $\alpha_{(k+1)\epsilon}(C) < 4n\sqrt{3n\rho}$. This proves part (i), which implies part (ii).

Remark 10. The theorem can be readily extended to the case of rectangular CV matrices as well as CV^T matrices $C = (\frac{1}{f\omega^i - s_j})_{i,j=0}^{m-1,n-1}$ for $|f| = 1$.

Next, for $((2k-1)\epsilon)$-approximate multiplication by a vector of ϵ-selectively ρ-neutered CV matrices C we yield nearly linear upper bounds on $\alpha_{(2k-1)\epsilon}(C)$, $\alpha_{(2k-1)\epsilon}(C^T)$, $\beta_{(2k-1)\epsilon}(C)$, and $\beta_{(2k-1)\epsilon}(C^T)$ provided that $\log(1/\epsilon) = O(\log(n))$.

Theorem 14. *Assume an ϵ-selectively ρ-neutered $n \times n$ CV matrix C for a valid 5-tuple $(h, k, n, \rho, \epsilon)$. Then it holds that (i) $\alpha_{(2k-1)\epsilon}(C) \leq 8n\rho + 4k\rho^2 + 6hn - n$, and so $\alpha_{(2k-1)\epsilon}(C) = 22n\rho - n + 5\delta_h n = O(n\rho)$ for $k = n/h$, $\rho = \lfloor h/2 \rfloor$ (cf. Remark 8), and $\delta_h = \lfloor h/2 - \rho \rfloor$, and therefore $\alpha_{(2k-1)\epsilon}(C) = O(n \log(n/\sqrt{\epsilon}))$ if $\rho = O(\log(n/\sqrt{\epsilon}))$ (cf. Remark 4). (ii) Furthermore if the matrix C is nonsingular and if $(2k-1)\epsilon||C^{-1}|| < 1/3$, then $\beta_{(2k-1)\epsilon}(C) = O(n\rho^2)$, and so $\beta_{(2k-1)\epsilon}(C) = O(n(\log(n/\sqrt{\epsilon}))^2)$ if $\rho = O(\log(n/\sqrt{\epsilon}))$.*

Proof. Combine Corollary 4 with part (i) of Theorem 9 for $l = u = \rho$, note that in this application of the theorem it holds that $\sum_{p=0}^{k-1} m_p = 3n$, $\sum_{q=0}^{k-1} n_q = n$, and $\sum_{q=0}^{k-1} m_q n_q \leq 3hn$, deduce that $\alpha_{(2k-1)\epsilon}(C) \leq 8n\rho + 4k\rho^2 + 6hn - n$, and obtain the bounds of part (i) of Theorem 14. Likewise, to deduce its part (ii) combine Corollary 4 with Theorem 4 and part (ii) of Theorem 9.

Theorem 15. *Suppose equation (10) expresses an $n \times n$ Vandermonde matrix $V_{\mathbf{s}}$ through the CV matrix $C = C_{\mathbf{s},f}$ for a complex f such that $|f| = 1$. Then (i) for any vector $\mathbf{u}$ one can approximate the vector $V_{\mathbf{s}}\mathbf{u}$ within the error norm bound $(\epsilon\sqrt{n})||\mathbf{u}|| \max_{i=0}^{n-1} |s_i^n - f|$ by using less than $4n\sqrt{7.2 \log_2(n/\sqrt{\epsilon})}$ flops. Furthermore suppose the CV matrix C is ϵ-selectively ρ-neutered for a valid 5-tuple $(h, k, n, \rho, \epsilon)$. Then (ii) one can approximate the product $V_{\mathbf{s}}\mathbf{u}$ within the error norm bound $((2k-1)\epsilon/\sqrt{n})||\mathbf{u}|| \max_{i=0}^{n-1} |s_i^n - f|$ by using $O(n \log(n/\sqrt{\epsilon}))$ flops and (iii) if the matrix $V_{\mathbf{s}}$ is nonsingular and if $(2k-1)\epsilon\sqrt{n}\mu_f||V_{\mathbf{s}}^{-1}|| < 1/3$, then one can approximate the vector $V_{\mathbf{s}}^{-1}\mathbf{u}$ within the error norm bound $2(2k-1)\epsilon\sqrt{n}\mu_f||V_{\mathbf{s}}^{-1}||\ ||\mathbf{u}||$ by using $O(n(\log(n/\sqrt{\epsilon})^2))$ flops where $|f| = 1$, $\mu_f = \max_{i=0}^{n-1} |s_i^n - f^n|^{-1}$, and one can choose f such that $\mu_f \leq n/\pi$. The same estimates hold where the transposes $V_{\mathbf{s}}^T$ and $-C_{f,\mathbf{s}} = C_{\mathbf{s},f}^T$ replace the matrices $V_{\mathbf{s}}$ and $C = C_{\mathbf{s},f}$, respectively, as well as for the problems of approximate evaluation of a polynomial of degree $n-1$ at the n knots $s_0, \dots, s_{n-1}$ and approximate interpolation to this polynomial from its n values at these knots.*

Proof. Deduce the claims about the matrix $V_{\mathbf{s}}$ and its transpose by combining Theorems 13 and 14 with equations (15)–(18) and note that μ_f is minimized under $|f| = 1$ for the complex points s_i equally spaced on the unit cirle $\{x : |x| = 1\}$. Extend the estimates to the case of polynomials by applying equation (5).

Remark 11. The proofs of Theorems 13–15 are constructive. They enable one to devise supporting algorithms that involve the computation of the centers c_q and c'_q (which define the basic neutered block columns N_q and N'_q) and subsequent computation of the ULV factorizations of the neutered blocks $\cap_q$ and $\cap'_q$ for all q. One can avoid a large part of this tedious computation and dramatically simplify the implementation, however, by following the papers [CGS07], [X12], [XXG12], and [XXCBa]: they bypass the computation of the centers c_q and c'_q and compute the HSS generators for the blocks defined by HSS trees instead of the generators for all neutered block columns. The algorithms of [CGS07], [X12], [XXCBa], and [XXG12] have been devised for a particular CV matrix $C_{1,\omega_{2n}}$, but can be readily extended to general CV and CV^T matrices.

Remark 12. In the extensive tests reported in [XXG12] the numerical rank of the matrix $C_{1,\omega_{2n}}$ consistently grew much slower than $\log(n)$ as n grew large. This empirical observation suggests that for a large class of CV matrices or possibly even for a "typical" CV matrix our upper estimate $\rho \le 1.2 \log_2(kn/\epsilon)$ of Remark 4 is overly pessimistic. Here are some more formal reasons for this conjecture. (i) Recall that the entries of a neutered block decay as block moves away from the diagonal blocks, which eventually decreases its ϵ-rank, but we have not used this observation in our analysis. (ii) For most of the dispositions of the set $\mathcal{S} = \{s_0, \dots, s_{n-1}\}$ on the complex plane and for most of the pairs i and j, our basic lower bound π/k on the values $|s_i - c_j|$ is overly pessimistic. (iii) We stated our estimate for ρ in Remark 4 to deduce our upper bound of (21) on the norm $||E||$, but this bound is not always sharp and is not always required. In particular, under the spectral matrix norm, our bounds for the errors of the approximation of the vectors $C\mathbf{u}$ and $C^{-1}\mathbf{u}$ are too generous and can be divided by $\sqrt{n}$. (iv) Likewise in the proof of Theorem 12 we deduced the bound $||(E_0, \dots, E_{2k-1})|| \le (2k-1)\epsilon$ from the bounds $||E_q|| \le \epsilon$ for $0, \dots, q = 2k - 1$ but the factor of $2k - 1$ would disappear from this bound if we used the column matrix norm $||\cdot||_1$. (v) We deduced our separation bound θ (in terms of the angles ϕ_i in the representation of the knots s_i in polar coordinates) for the worst case absolute values $|s_i|$, but we can expect that the bound is substantially stronger in the average case.

8 Extension of matrix structures and the cost estimates

To extend our progress we recall and employ the Sylvester displacements $AM - MB$ of the matrices with the structures of Cauchy and Vandermonde types. We say that an $n \times n$ matrix $M = C_{\mathbf{s},\mathbf{t}}(F, G)$ has a $(\mathbf{s}, \mathbf{t}, d)$-structure of the Cauchy type $(\mathbf{s}, \mathbf{t})$ if $D_{\mathbf{s}}M - MD_{\mathbf{t}} = FG^T$ for two vectors $\mathbf{s}$ and $\mathbf{t}$, the diagonal matrices $D_{\mathbf{s}}$ and $D_{\mathbf{t}}$, and $n \times d$ matrices F and G having columns $\mathbf{f}_j$ and $\mathbf{g}_j$, for $j = 0, \ldots, d-1$, and rows $\mathbf{u}_i$ and $\mathbf{v}_i$, for $i = 0, \ldots, n-1$, respectively, or equivalently if

$$M = \sum_{j=1}^{d} D(\mathbf{f}_j) C D(\mathbf{g}_j) = \left(\frac{\mathbf{u}_i^T \mathbf{v}_j}{s_i - t_j} \right)_{i,j=0}^{n-1} \text{ for } C = \left(\frac{1}{s_i - t_j} \right)_{i,j=0}^{n-1} \tag{25}$$

(cf. [P01, Examples 1.4.1 and 4.6.4]). In this case we write $M = C_{\mathbf{s},\mathbf{t}}(F, G)$. In particular $M = (\frac{1}{s_i - t_j})_{i,j=0}^{n-1}$ and $d = 1$ if $F = G = (1, \ldots, 1)^T$. Formula (25) expresses a matrix M via a nonunique generating pair $\{F, G\}$ of length d for its displacement $D_{\mathbf{s}} M - M D_{\mathbf{t}}$, and this pair is called a *displacement generator* of the matrix M. For small integers d we call the matrices M of (25) *Cauchy-like* and having *structure of Cauchy type*. Their multiplication by a vector is reduced to d multiplications of the Cauchy matrix C by d vectors and $O(n)$ other flops.

More generally suppose that the displacement $AM - MB$ of a matrix M has a small rank (called *displacement rank* of the matrix M) for some pair of *operator matrices* A and B. Then we say that the matrix M has the structure defined by the pair (A, B). For example, the above pair $(A, B) = (D_{\mathbf{s}}, D_{\mathbf{t}})$ defines the structure of Cauchy type. For another example, let $Z_c = \begin{pmatrix} \mathbf{0}^T & c \\ I_{n-1} & \mathbf{0} \end{pmatrix}$ for a scalar c denote the $n \times n$ matrix of c-circular shift and let c and d be two distinct constants. Then the pairs $(A, B) = (Z_c, Z_d)$ and $(A, B) = (Z_c, Z_d^T)$ define the classes of matrices with Toeplitz and Hankel structures, respectively, whereas the pair $(A, B) = (D_{\mathbf{s}}, Z_c)$ of operator matrices for any complex scalar c defines an extension of the class of the Vandermonde matrices $V = V_{\mathbf{s}} = (s_i^{j-1})_{i,j=0}^{n-1}$, and it holds that

$$D_{\mathbf{s}} V - V Z_c = (s_i^n - c)_{i=0}^{n-1} (0 \mid \ \ldots \ \mid 0 \mid 1). \tag{26}$$

By using the transpose of this pair $(Z_c^T, D_{\mathbf{s}})$ we extend the class of transposed Vandermonde matrices V^T, and it holds that

$$Z_c^T V^T - V^T D_{\mathbf{s}} = (0 \mid \ \ldots \ \mid 0 \mid 1)^T ((s_i^n - c)_{i=0}^{n-1})^T. \tag{27}$$

We refer the reader to [BP94], [GO94], [P01, Sections 4.1–4.5], [PW03], and the bibliography therein on expressing various structured matrices via their displacements (including matrices with the structures of Toeplitz, Hankel, Cauchy, Vandermonde and transposed Vandermonde types) and to [P01, Section 1.5] on performing arithmetic operations with matrices in terms of their displacements, exemplified by the following simple but basic result.

Theorem 16. *(Cf. [P90] and [P01, Theorem 1.5.4].) Assume five matrices A, B, C, M and N with compatible sizes. Then*

$$A(MN) - (MN)C = (AM - MB)N + M(BN - NC), \tag{28}$$

and so if $AM - MB = G_M H_M^T$ and $BN - NC = G_N H_N^T$, then $A(MN) - (MN)C = G_{MN} H_{MN}^T$ for $G_{MN} = (G_M \mid M G_N)$ and $H_{MN} = (N^T H_M \mid H_N)$.

By virtue of this theorem one can transform a pair of operator matrices (A, B) into any other pair of operator matrices by means of multiplication with appropriate multipliers. Consequently one can transform the classes of matrices with the structures of Toeplitz, Hankel, Vandermonde, transposed Vandermonde, and Cauchy types into each other at will. This fact has broad algorithmic impact, because it enables one to exploit distinct features of various matrix structures, in particular (as we noted in Section 4) the invariance of the matrix structure of Cauchy type in row and column interchange (in contrast to the structures of Toeplitz and Hankel types), exploited in [GKO95] and [G98], and the expansion of the matrix entries into Loran's series (unlike the structures of the three other types), exploited in [MRT05], [CGS07], and [XXG12].

Remark 13. Transformation of the matrix structures from Vandermonde type into Cauchy type was the initial step for our current progress (cf. (15)–(18)). Transformation into the opposite direction seems to be also highly promising. Namely exact or approximate polynomial evaluation and interpolation (cf. Problems 1 and 2) are equivalent to computing or approximating the products of a Vandermonde matrix and its inverse by a vector. Now assume any successful algorithms for these tasks and combine equation (9) and Remark 1 to extend them to multiplication of a Cauchy matrix and its inverse by a vector and consequently to Problems 3 and 4 as well as some other problems of rational interpolation [P01] and a broad area of computations with matrices having structures of Vandermonde and Cauchy types [Pa].

9 Conclusions

In this paper we employ transformations of matrix structures and some variants of the Multipole techniques to accelerate dramatically the known approximation algorithms for the nth degree polynomial interpolation and multipoint evaluation. The basic computational blocks of our algorithms are the FFT and Multipole type algorithms, well studied, efficiently implemented, and showing sufficiently stable numerical behavior in practice.

Our progress relies on our reduction of polynomial and Vandermonde matrix computations to HSS matrix computations. As soon as we achieve this reduction, we just invoke the known efficient algorithms of [MRT05], [CGS07], and [XXG12]. Overall this accelerates ϵ-approximate multipoint polynomial evaluation by a factor of $\sqrt{n/\log n}$ as long as $\log(1/\epsilon) = O(\log(n))$, and furthermore for a large class of inputs we support nearly optimal arithmetic cost bounds for both ϵ-approximate evaluation and interpolation. A natural research subject is the study of this class, which we specify by associating it with ϵ-selectively ρ-neutered CV matrices for valid 5-tuples $(h, k, n, \rho, \epsilon)$ (cf. Definition 5).

As a potential beneficiary of our progress we recall the classical task of univariate polynomial root-finding. The current best package of subroutines MPSolve [BF00] relies on Ehrlich–Aberth iterations, which amount essentially to recursive multipoint polynomial evaluation. MPSolve performs it in quadratic time by means of Horner's algorithm, versus nearly linear time of our algorithms. For a large class of inputs a number of the initial iterations can go with the IEEE standard double precision, and at these stages our algorithms have clear advantage over Horner's.

Transformations of matrix structures can extend our progress to computations with various matrices having structures of Vandermonde and Cauchy types, for example, confluent Vandermonde matrices and Loewner matrices, and various problems of rational interpolation such as the Nevanlinna–Pick and matrix Nehari problems, although guarding against numerical problems may limit these extensions.

Other natural subjects for further formal and experimental study, include the ones pointed out in Remarks 7 and 12, and there we can add the computation of least-squares solutions of Vandermonde, CV and CV^T linear systems of equations (cf. [XXCBa]) as well as the estimation of the threshold input sizes for which our algorithms, running in nearly linear time, outperform the known numerical algorithms, running in quadratic time, and the algorithms supporting Theorem 13 and part (i) of Theorem 15. Finally, one can facilitate parallel implementation of our algorithms by using the known efficient parallel algorithms for FFT and nested dissection (see, e.g., [B99], [GS66], [P93], and [PR93]).

Figure 1. The block diagonal is shown in black color.

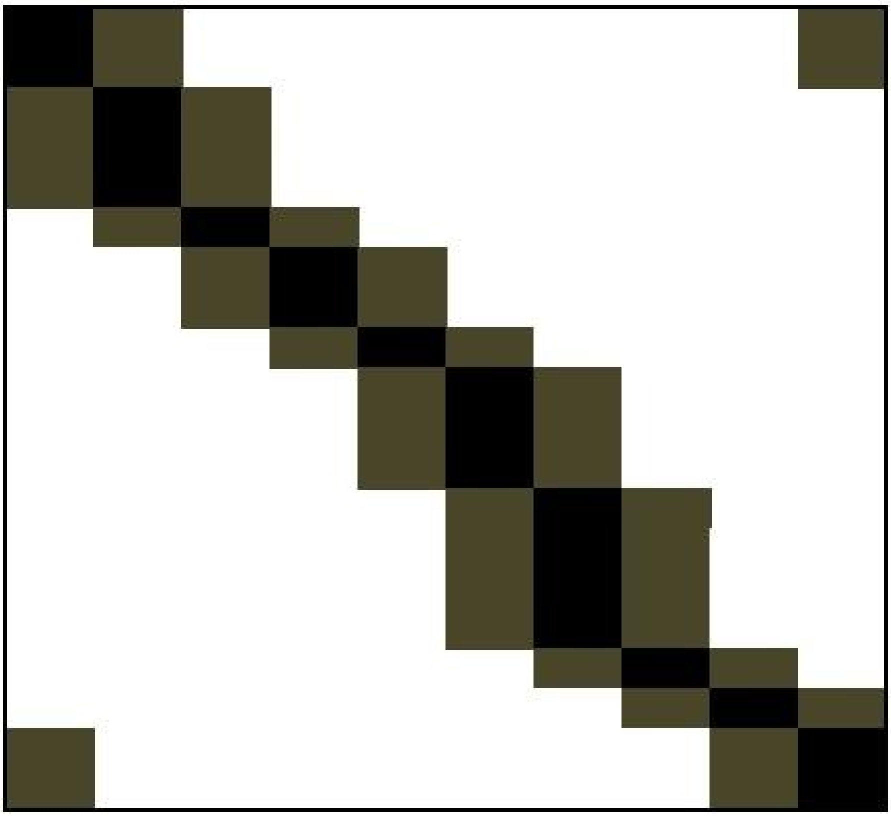

Figure 2. Two basic neutered block columns are shown in blue and pink.

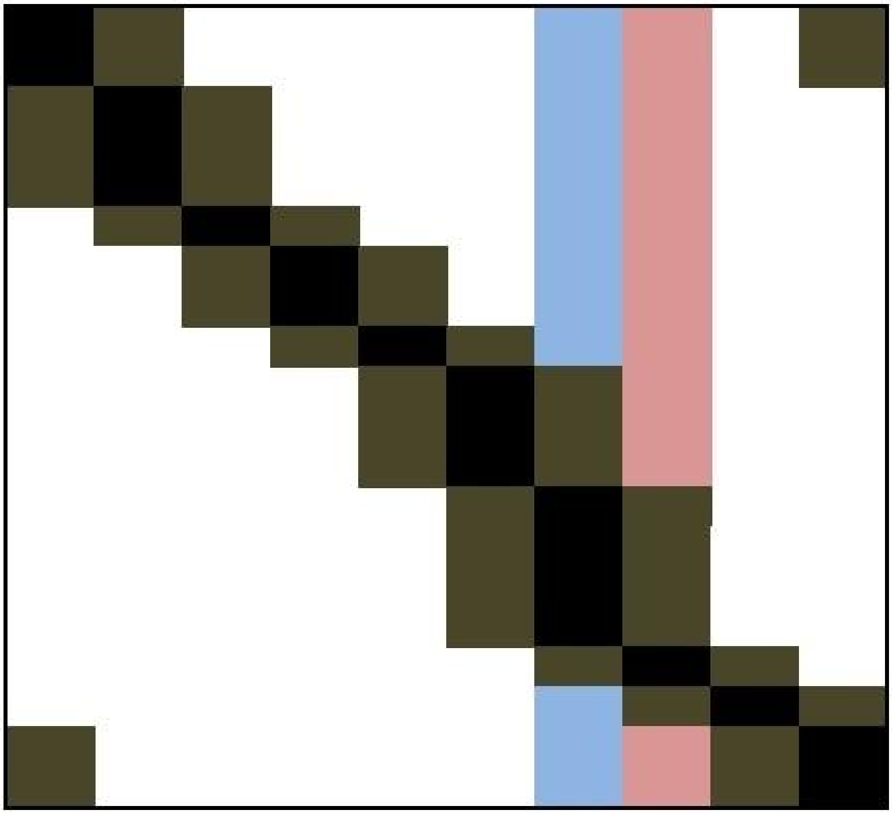

Figure 3. The neutered union of two basic neutered block columns is shown in green. The rest of them is shown in blue and pink.

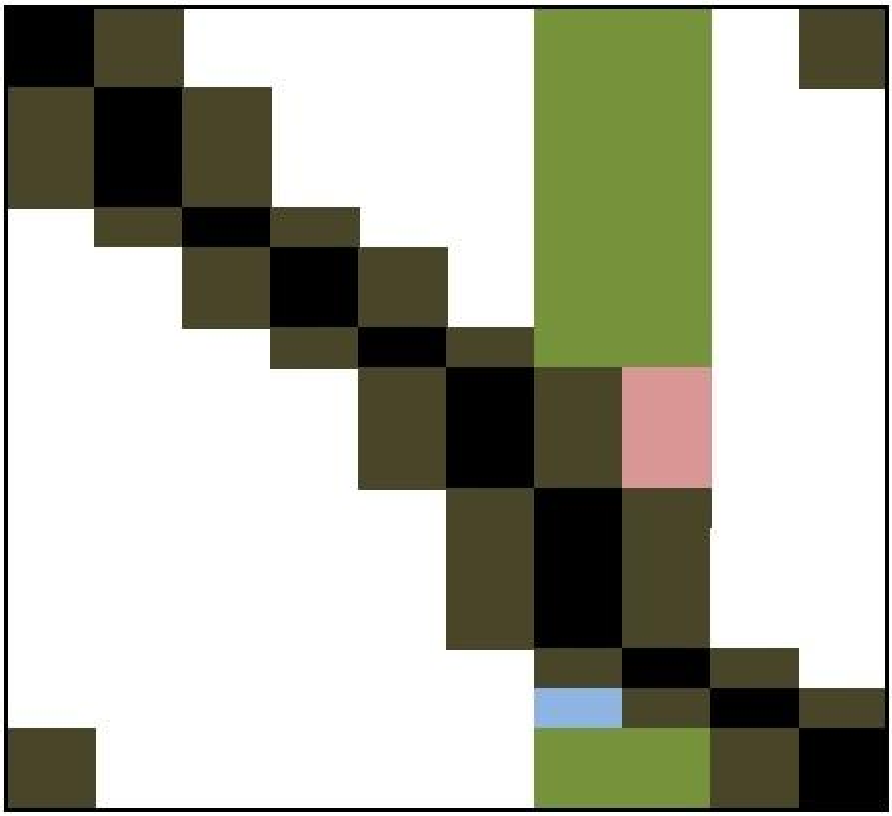

Figure 4. Blue lines bound the sectors Γ_0, Γ_1, Γ_3, Γ_{k-2}, and Γ_{k-1}.

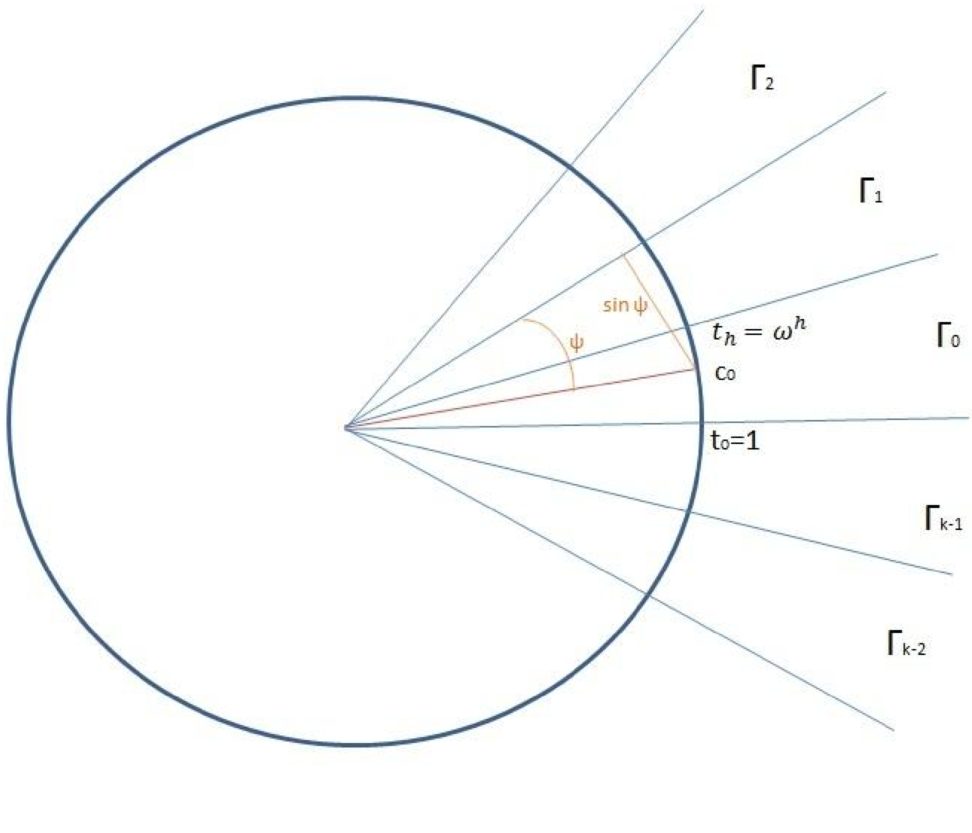

Figure 5. Blue lines bound the sectors Γ_0, Γ_1, Γ_3, Γ_{k-2}, and Γ_{k-1}. Purple lines bound the sectors Γ_0, Γ_1, Γ_3, Γ_{k-2}, and Γ_{k-1}.

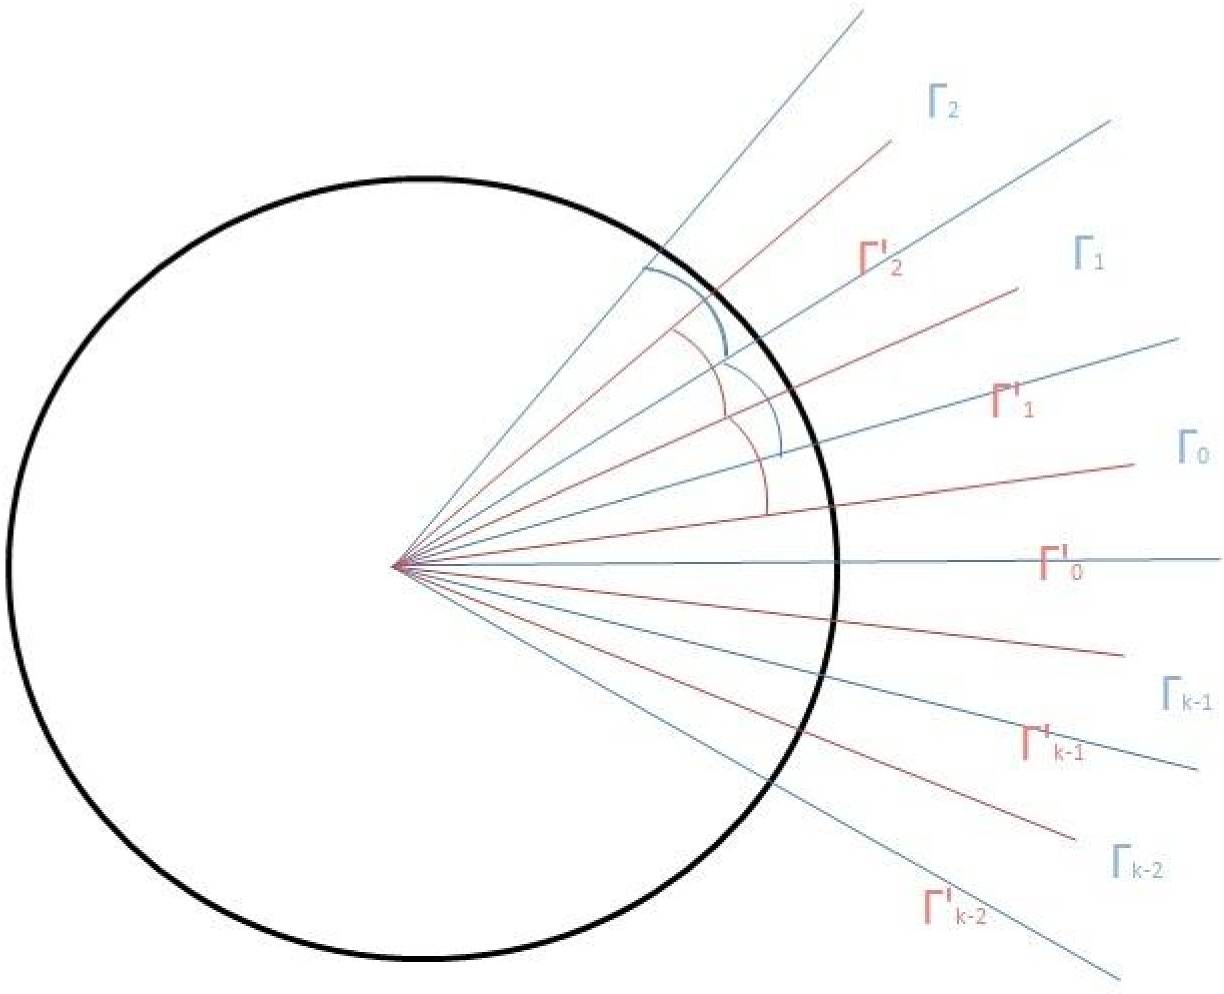

Figure 6. A basic neutered block column N_q is shown in green. The remaining parts of two basic neutered block columns N_q and N_{q+1} are shown in blue and pink, respectively.

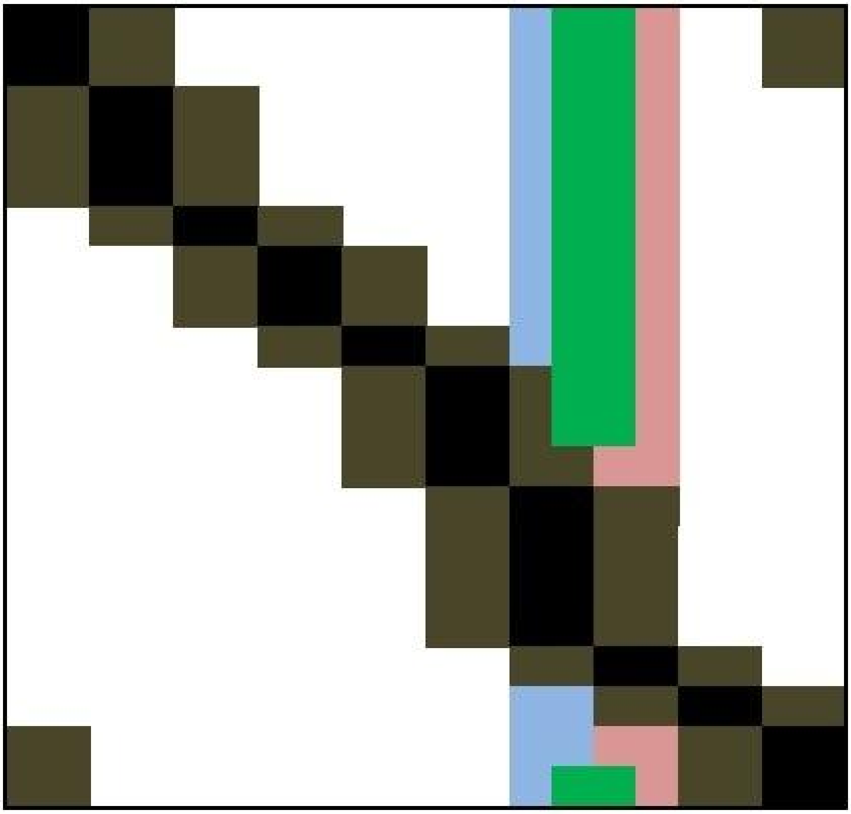

Acknowledgements: Our research has been supported by the NSF Grant CC 1116736 and the PSC CUNY Awards 64512–0042 and 65792–0043.

References

[B99] R. Bracewell, *The Fourier Transform and Its Applications*. McGraw-Hill, New York, 1999 (3rd edition).

[BEGO08] T. Bella, Y. Eidelman, I. Gohberg, V. Olshevsky, Computations with Quasiseparable Polynomials and Matrices, *Theoretical Computer Science, Special Issue on Symbolic–Numerical Algorithms* (D. A. Bini, V. Y. Pan, and J. Verschelde editors), **409**, **2**, 158–179, 2008.

[BF00] D. A. Bini, G. Fiorentino, Design, Analysis, and Implementation of a Multiprecision Polynomial Rootfinder, *Numer. Algs.*, **23**, 127–173, 2000.

[B-O83] M. Ben-Or, Lower Bounds for Algebraic Computation Trees, *Proceedings of 15th Annual ACM Symposium on Theory of Computing (STOC'83)*, 80–86, ACM Press, New York, 1983.

[BP70] A. Björck, V. Pereyra, Solution of Vandermonde Systems of Equations, *Math. of Computation*, **24**, 893–903, 1970.

[BP94] D. Bini, V. Y. Pan, *Polynomial and Matrix Computations, Volume 1: Fundamental Algorithms*, Birkhäuser, Boston, 1994.

[CGR98] J. Carrier, L. Greengard, V. Rokhlin, A Fast Adaptive Algorithm for Particle Simulation, *SIAM J. Scientific Computing*, **9**, 669–686, 1998.

[CGS07] S. Chandrasekaran, M. Gu, X. Sun, J. Xia, J. Zhu, A superfast algorithm for Toeplitz systems of linear equations, *SIAM J. Matrix Anal. Appl.*, **29**, 1247–1266, 2007.

[EG02] Y. Eidelman, I. Gohberg, A Modification of the Dewilde–van der Veen Method for Inversion of Finite Structured Matrices, *Linear Algebra and Its Applications*, **343**, 419–450, 2002.

[EGHa] Y. Eidelman, I. Gohberg, I. Haimovici, *Separable type representations of matrices and fast algorithms. Volume 1. Basics. Completion problems. Multiplication and inversion algorithms*, Birkhauser, 2013.

[EGHb] Y. Eidelman, I. Gohberg, I. Haimovici, *Separable type representations of matrices and fast algorithms. Volume 2. Eigenvalue method*, Birkhauser, 2013.

[F72] C. M. Fiduccia, Polynomial Evaluation via the Division Algorithm: The Fast Fourier Transform Revisited, *Proc. 4th Annual ACM Symp. on Theory of Computing (STOC'72)*, 88–93, 1972.

[G98] M. Gu, Stable and Efficient Algorithms for Structured Systems of Linear Equations, *SIAM J. Matrix Anal. Appl.*, **19**, 279–306, 1998.

[GI88] W. Gautschi, G. Inglese, Lower Bounds for the Condition Number of Vandermonde Matrices, *Numerische Mathematik*, **52**, 241–250, 1988.

[GG03] J. von zur Gathen, J. Gerhard, *Modern Computer Algebra*, Cambridge University Press, Cambridge, UK, 2003 (second edition).

[GGS87] A. Gerasoulis, M. D. Grigoriadis, L. Sun, A Fast Algorithm for Trummer's Problem, *SIAM Journal on Scientific and Statistical Computing*, **8, 1**, 135–138, 1987.

[GKO95] I. Gohberg, T. Kailath, V. Olshevsky, Fast Gaussian Elimination with Partial Pivoting for Matrices with Displacement Structure, *Mathematics of Computation*, **64**, 1557–1576, 1995.

[GL96] G. H. Golub, C. F. Van Loan, *Matrix Computations*, 3rd edition, The Johns Hopkins University Press, Baltimore, Maryland, 1996.

[GO94] I. Gohberg, V. Olshevsky, Complexity of Multiplication with Vectors for Structured Matrices, *Linear Algebra Appls.*, **202**, 163–192, 1994.

[GR87] L. Greengard, V. Rokhlin, A Fast Algorithm for Particle Simulation, *Journal of Computational Physics*, **73**, 325–348, 1987.

[GS66] W. Gentelman, G. Sande, Fast Fourier Transform for Fun and Profit, *Full Joint Comput. Conference*, **29**, 563–578, 1966.

[H72] E. Horowitz, A Fast Method for Interpolation Using Preconditioning, *Information Processing Letters*, **1, 4**, 157–163, 1972.

[KZ08] S. Köhler, M. Ziegler, On the Stability of Fast Polynomial Arithmetic, *Proc. 8-th Conference on Real Numbers and Computers* (J.D. Bruguera, M. Daumas, eds.), 147–156, 2008.

[LRT79] R. J. Lipton, D. Rose, R. E. Tarjan, Generalized Nested Dissection, *SIAM J. on Numerical Analysis*, **16**, **2**, 346–358, 1979.

[MB72] R. Moenck, A. Borodin, Fast Modular Transform via Division, *Proceedings of 13th Annual Symposium on Switching and Automata Theory*, 90–96, IEEE Computer Society Press, Washington, DC, 1972.

[MRT05] P. G. Martinsson, V. Rokhlin, M. Tygert, A fast algorithm for the inversion of Toeplitz matrices, *Comput. Math. Appl.*, **50**, 741–752, 2005.

[P64] F. Parker, Inverses of Vandermonde matrices, *Amer. Math. Monthly*, **71**, 410–411, 1964.

[P66] V. Y. Pan, On Methods of Computing the Values of Polynomials, *Uspekhi Matematicheskikh Nauk*, **21**, **1(127)**, 103–134, 1966. [Transl. *Russian Mathematical Surveys*, **21, 1(127)**, 105–137, 1966.]

[P90] V. Y. Pan, On Computations with Dense Structured Matrices, *Math. of Computation*, **55, 191**, 179–190, 1990. Proceedings version in *Proc. Intern. Symposium on Symbolic and Algebraic Computation (ISSAC'89)*, 34–42, ACM Press, New York, 1989.
5

[P91]

[P93] V. Y. Pan, Parallel Solution of Sparse Linear and Path Systems, in *Synthesis of Parallel Algorithms* (J.H. Reif, editor), Chapter 14, pp. 621–678, Morgan Kaufmann publishers, San Mateo, California (1993).

[P95] V. Y. Pan, An Algebraic Approach to Approximate Evaluation of a Polynomial on a Set of Real Points, *Advances in Computational Mathematics*, **3**, 41–58, 1995.

[P01] V. Y. Pan, *Structured Matrices and Polynomials: Unified Superfast Algorithms*, Birkhäuser/Springer, Boston/New York, 2001.

[Pa] V. Y. Pan, Transformations of Matrix Structures Work Again, Tech. Report TR 2013004, *PhD Program in Comp. Sci., Graduate Center, CUNY*, 2013
Available at http://www.cs.gc.cuny.edu/tr/techreport.php?id=446

[PR93] V. Y. Pan, J. Reif, Fast and Efficient Parallel Solution of Sparse Linear Systems, *SIAM J. on Computing*, **22**, **6**, 1227–1250, 1993.

[PSLT93] V. Y. Pan, A. Sadikou, E. Landowne, O. Tiga, A New Approach to Fast Polynomial Interpolation and Multipoint Evaluation, *Computers and Math. (with Applications)*, **25, 9**, 25–30, 1993.

[PW03] V. Y. Pan, X. Wang, Inversion of Displacement Operators, *SIAM J. on Matrix Analysis and Applications*, **24, 3**, 660–677, 2003.

[PZHY97] V. Y. Pan, A. Zheng, X. Huang, Y.Yu, Fast Multipoint Polynomial Evaluation and Interpolation via Computation with Structured Matrices, *Annals of Numerical Math.*, **4**, 483–510, 1997.

[R85] V. Rokhlin, Rapid Solution of Integral Equations of Classical Potential Theory, *Journal of Computational Physics*, **60**, 187–207, 1985.

[S73] V. Strassen, Die Berechnungskomplexetät von elementarysymmetrischen Funktionen und von Interpolationskoeffizienten, *Numerische Mathematik*, **20, 3**, 238–251, 1973.

[S98] G. W. Stewart, *Matrix Algorithms, Vol I: Basic Decompositions*, SIAM, Philadelphia, 1998.

[VVGM05] R. Vandebril, M. Van Barel, G. Golub, N. Mastronardi, A Bibliography on Semiseparable Matrices, *Calcolo*, **42, 3–4,** 249–270, 2005.

[VVM07] R. Vandebril, M. Van Barel, N. Mastronardi, *Matrix Computations and Semiseparable Matrices: Linear Systems* (Volume 1), The Johns Hopkins University Press, Baltimore, Maryland, 2007.

[VVM08] R. Vandebril, M. Van Barel, N. Mastronardi, *Matrix Computations and Semiseparable Matrices: Eigenvalue and Singular Value Methods* (Volume 2), The Johns Hopkins University Press, Baltimore, Maryland, 2008.

[X12] J. Xia, On the Complexity of Some Hierarchical Structured Matrix Algorithms, *SIAM J. Matrix Anal. Appl.*, **33**, 388–410, 2012.

[XXCBa] J. Xia, Y. Xi, S. Cauley, and V. Balakrishnan, Superfast and stable structured solvers for Toeplitz least squares via randomized sampling, *SIAM J. Matrix Anal. and Applications*, in press.

[XXG12] J. Xia, Y. Xi, M. Gu, A superfast structured solver for Toeplitz linear systems via randomized sampling, *SIAM J. Matrix Anal. Appl.*, **33**, 837–858, 2012.