BIT manuscript No.
(will be inserted by the editor)

Computation of connection coefficients and measure modifications for orthogonal polynomials

Akil Narayan · **Jan S. Hesthaven**

Received: date / Accepted: date

Abstract We observe that polynomial measure modifications for families of univariate orthogonal polynomials imply sparse connection coefficient relations. We therefore propose connecting L^2 expansion coefficients between a polynomial family and a modified family by a sparse transformation. Accuracy and conditioning of the connection and its inverse are explored. The connection and recurrence coefficients can simultaneously be obtained as the Cholesky decomposition of a matrix polynomial involving the Jacobi matrix; this property extends to continuous, non-polynomial measure modifications on finite intervals. We conclude with an example of a useful application to families of Jacobi polynomials with parameters (γ, δ) where the fast Fourier transform may be applied in order to obtain expansion coefficients whenever 2γ and 2δ are odd integers.

Keywords orthogonal polynomials · measure modifications · connection coefficients · Jacobi polynomials · fast Fourier transform

Mathematics Subject Classification (2000) 65 · 41

1 Introduction

Orthogonal polynomials are used in many applications for their attractive approximation properties. On an interval $I \subset \mathbb{R}$, a given positive measure $\mathrm{d}\omega$ with Lebesgue-Radon-Nikodym derivative ω results in a well-defined family $\{\pi_n\}_{n=0}^{\infty}$ of univariate L^2_ω-orthogonal polynomials. For a measurable semi-positive function $q : I \to \mathbb{R}$, the so-called modification problem is the determination of relevant quantities for the

Akil Narayan
Department of Mathematics Purdue University 150 N. University Street, West Lafayette, IN 47907
Tel.: +1-765-4946714
Fax.: +1-765-4940548
E-mail: akil@purdue.edu

Jan S. Hesthaven
Division of Applied Mathematics Brown University, Box F, 182 George Street, Providence, RI 02912

modified orthogonal polynomial family p associated with the weight function $q\omega$. The modification problem is a classical problem, and the literature boasts deep understanding and fruitful results for practical computation. However, most of these techniques concentrate on the determination of the three-term recurrence coefficients (equivalently, the Jacobi matrix) associated with the modified family p (see [6,8, 15]). We explore a slightly different avenue: assuming q is a polynomial of degree K, a result due to Christoffel states that qp_n is a linear combination of $K+1$ polynomials π_m. Our goal is the determination of the coefficients that define this relation; these coefficients allow us to translate L^2 expansion coefficients with a relation that mirrors Christoffel's result, thus enabling a *sparse* coefficient connection between the families. For N degrees of freedom, the sparse connection result for polynomial modifications can be used to effect an $\mathcal{O}(N)$-cost transform between expansion coefficients for π and those for p. This result can be extended to show that if J is the Jacobi matrix for π, then a leading principal submatrix of $q(J)$ is positive and the Cholesky decomposition of this matrix gives exactly the sought connection coefficients. This is a variant of the classical Kautsky/Golub Jacobi matrix characterization of $q(J)$ [19] – we extend this to cases when q is a non-polynomial.

In Section 2 we introduce notation and known results on orthogonal polynomials necessary for our discussion. Section 3 defines the resulting approximation operators, explores conditioning of the procedures, and establishes the Cholesky decomposition result. We discuss two computational techniques for computing the connection coefficients in Section 4. Finally, Section 5 validates the speed and accuracy of the method; in particular, we present the straightforward application to polynomial modifications of the Chebyshev family, resulting in a competitive fast discrete polynomial transform for Jacobi polynomials of family (γ,δ) when 2γ and 2δ are odd integers.

2 Orthogonal polynomials

We review some basic properties of orthogonal polynomials; the books [26] and [11] are excellent references for this subject. Let the interval of approximation $I \subset \mathbb{R}$ be an open interval and denote $\bar{I}$ its closure on the extended real line. $C(\bar{I})$ denotes the collection of functions continuous on $\bar{I}$. Let $\omega : I \to \mathbb{R}^+$ be a weight function associated with a measure that is semi-positive on I; we assume that $\mathrm{d}\omega$ has finite polynomial moments of all orders: $\int_I r^n \mathrm{d}\omega < \infty$. For $p \geq 1$ let L^p_ω denote the space of real-valued functions on I whose magnitude p-th power is integrable. We will concentrate on $p = 2$ with inner product and induced norm $\langle f, f, \rangle_\omega = \|f\|^2_\omega = \int_I |f^2| \mathrm{d}\omega$. The polynomials that are L^2_ω orthonormal against the weight ω are denoted $\pi_n[\omega] = \pi_n$ and the derivatives $\pi_n^{(d)}$ satisfy the standard four-term recurrence:

$$\sqrt{\beta_{n+1}}\pi^{(d)}_{n+1} = (r-\alpha_n)\pi^{(d)}_n - \sqrt{\beta_n}\pi^{(d)}_{n-1} + d\pi^{(d-1)}_n, \tag{2.1}$$

for $d,n \geq 0$ where $\pi_n^{(d)} \equiv 0$ for $n < d, n < 0$ and $\pi_d^{(d)} = d! \prod_{j=0}^d 1/\sqrt{\beta_j}$. The recurrence coefficients α, β are known for many classical families of polynomials; the above indicates that knowledge of the recurrence coefficients α_n, β_n allows evaluation of derivatives of all orders. By (2.1), π_n is the unique polynomial with positive

leading coefficient. Orthonormality implies $\langle \pi_n, \pi_m \rangle_\omega = \delta_{m,n}$, where $\delta_{m,n}$ is the Kronecker delta symbol. For any $N > 0$, the $N \times N$ symmetric tridiagonal Jacobi matrix for this orthogonal polynomial family is defined in terms of the recurrence coefficients:

$$J_N^\pi = \begin{pmatrix} \alpha_0 & \sqrt{\beta_1} & 0 & \cdots & \\ \sqrt{\beta_1} & \alpha_1 & \sqrt{\beta_2} & 0 & \cdots \\ 0 & \ddots & \ddots & \ddots & \\ & & & \sqrt{\beta_{N-1}} & \alpha_{N-1} \end{pmatrix}.$$

Orthogonal polynomial families satisfy the Christoffel-Darboux identity, given by

$$\sum_{j=0}^{n-1} \pi_j(x)\pi_j(y) = \sqrt{\beta}_n \frac{\pi_n(x)\pi_{n-1}(y) - \pi_n(y)\pi_{n-1}(x)}{x-y}, \tag{2.2}$$

for all $n \geq 1$. When $r \notin I$, define the ratio of successive polynomial evaluations as $\tau_n(r) = \pi_n(r)/\pi_{n-1}(r)$ for $n > 0$ and $\tau_0 = \pi_0$. Then the following recurrence is easily derived from the $d = 0$ version of (2.1)

$$\sqrt{\beta_1}\tau_1 = (r - \alpha_0), \qquad \sqrt{\beta_{n+1}}\tau_{n+1} = (r - \alpha_n) - \frac{\sqrt{\beta_n}}{\tau_n}, \tag{2.3}$$

when $n > 0$. The above relation allows stable evaluation of polynomial ratios when overflow of $\pi_n(r)$ becomes an issue.

We denote polynomials and related quantities with monic normalization (unity leading coefficient) with overhead tilde's: i.e. $\tilde{\pi}_n(r)$ and $\tilde{\tau}_n = \tilde{\pi}_n/\tilde{\pi}_{n-1}$. The monic and orthonormal polynomials are related by $\pi_n = \kappa_n \tilde{\pi}_n$, with $\kappa_n = \prod_{j=0}^n 1/\sqrt{\beta_j}$ for all $n \geq 0$ the leading coefficient of π_n. The recurrence quantities α, β do not have tilde variations, and obviously the monic leading coefficient is $\tilde{\kappa}_n \equiv 1$, so it is never used.

Define the dimension-N space $\Pi_{N-1} = \operatorname{span}\{r^n : 0 \leq n \leq N-1\}$. For any $f \in L^2_\omega$, define the coefficient

$$\hat{f}_n^\pi = \int_I f\, \pi_n \, \mathrm{d}\omega = \langle f, \pi_n \rangle_\omega \tag{2.4}$$

and the resulting infinite sum $\sum_{n=0}^\infty \hat{f}_n^\pi \pi_n(r)$ is an L^2_ω representation for f. The superscripts on the coefficients indicate the polynomial family to which the coefficients correspond. An approximating truncated sum can be produced via a projection operator, which we define as

$$\mathscr{P}_N^\pi f = \sum_{n=0}^{N-1} \hat{f}_n^\pi \pi_n.$$

Then assuming completeness of polynomials in L^2_ω, we have

$$\begin{aligned} &\|f - \mathscr{P}_n^\pi f\|_\omega \longrightarrow 0, \;\; n \to \infty, \\ &\langle f - \mathscr{P}_n^\pi f, \phi \rangle_\omega = 0, \;\; \phi \in \Pi_{n-1} \end{aligned} \tag{2.5}$$

Let the N-point Gauss quadrature rule associated with the polynomial family π have nodes and weights $\{r_n^\pi\}_{n=1}^N$ and $\{w_n^\pi\}_{n=1}^N$, respectively. I.e.:

$$\int_I f \mathrm{d}\omega \simeq \sum_{n=1}^N f(r_n^\pi) w_n^\pi.$$

We assume the nodes are ordered and suppress the dependence of r_n^π and w_n^π on N. The Gaussian quadrature rule satisfies

$$\sum_{n=1}^N \phi(r_n^\pi) w_n^\pi = \int_I \phi \, \mathrm{d}\omega, \qquad \phi \in \Pi_{2N-1}. \tag{2.6}$$

The nodes r_n^π may be computed as eigenvalues of the Jacobi matrix. The first components of normalized eigenvectors gives the weights w_n^π (see e.g. [14]), or one may employ (2.2) and (2.6) to yield $w_n^\pi = (\sum_{j=0}^{N-1} \pi_j(r_n^\pi)^2)^{-1}$. The cost to compute the eigenvalues of J_N^π and to evaluate the weights for the quadrature rule is $\mathscr{O}(N^2)$. Notable exceptions are special Chebyshev-like families when explicit, simple expressions are known for the quadrature nodes and weights. However, asymptotic $\mathscr{O}(N)$ algorithms exist for any ω, at least for the determination of the nodes [13].

2.1 Connection Coefficients

The problem of rewriting a truncated expansion of one family of orthogonal polynomials $\pi[\omega]$ into another family $p[w]$ can be cast as the problem of determining the connection coefficients $\lambda_{n,m}^{p,\pi}$:

$$\pi_n = \sum_{m=0}^n p_m \lambda_{n,m}^{p,\pi}. \tag{2.7}$$

By orthogonality, the λ coefficients are given by

$$\lambda_{n,m}^{p,\pi} = \langle \pi_n, p_m \rangle_w. \tag{2.8}$$

If the recurrence coefficients for both families p and π are known, one may directly compute the connection coefficients [21], but we focus on the situation when we do not have knowledge of the recurrence coefficients for p. Furthermore, we concentrate on the case when the family p is a modification of the family π: this gives us knowledge of certain sparsity of the connection coefficients.

Let us assume that the λ coefficients can be computed; a worthwhile consideration is the process of transforming the truncated expression $\mathscr{P}_N^\pi f$ into $\mathscr{P}_N^p f$. The expansion coefficients from each family can be directly related from (2.7):

$$\hat{f}_n^p = \sum_{m=n}^\infty \hat{f}_m^\pi \lambda_{m,n}^{p,\pi}$$

Thus for an exact connection we require knowledge of an infinite number of π-expansion and λ connection coefficients. Since we are working with truncated expansions, we do not know all the coefficients $\hat{f}_n^\pi$ for $n \geq N$, and thus some error will

be present. Even simply evaluating the connection (2.7) for each $n = 0, 1, \ldots, N-1$ requires $\mathscr{O}(N^2)$ operations. In the following section we observe that when p and π are related in a certain way, the connection is 'sparse' (meaning most of the coefficients vanish) and the change-of-basis can be implemented in $\mathscr{O}(N)$ cost, and we can quantify the error made with knowledge of only finitely many expansion coefficients.

2.2 Modification Problems

For a general family of orthogonal polynomials $\pi[\omega]$ with known recurrence, a modification problem is the problem of computing information (e.g. the recurrence) of a new family $p[w]$, where the weight functions are related by $q\,\omega = w$ for some function q semi-positive on I. (Note that allowing q to become negative makes the resulting weighted integral violate positivity and thus it would not define a norm.) In this paper we focus on the particular case where $q \in \Pi_K$ for some $K \in \mathbb{N}_0$; that is, a polynomial modification. This problem is closely related to the connection problem of the previous section. The reason we restrict ourselves to a polynomial modification is because of a fundamental result of Christoffel's: with the weight functions related as described and $q \in \Pi_K$, then

$$qp_n = \sum_{k=0}^{K} C_{n,n+k}\pi_{n+k}, \tag{2.9}$$

where the $C_{n,n+k}$ are constants. To see that (2.9) is true, note that $\langle p_n, p_m\rangle_w = 0\ \forall m < n$ implies that $\langle qp_n, \pi_m\rangle_\omega = 0\ \forall m < n$. Explicit formulas for the constants $C_{n,n+k}$ were first explored by Christoffel and later Uvarov when q is a rational function (the latter case is not considered here). For a more complete discussion with historical references we refer to Section 2.5 in [26] or Section 2.4 in [11]. The utility we find in (2.9) is that it can be used to relate expansion coefficients; in particular, the modification coefficients $C_{n,n+k}$ and the connection coefficients $\lambda_{n,m}$ are one and the same.

Proposition 2.1 *Let weight functions ω and w be related by $w = q\omega$ with $q \in \Pi_K$; denote the associated polynomial families $\pi[\omega]$ and $p[w]$. Assume $f \in L^2_\omega$. Then*

$$\hat{f}^p_n = \sum_{m=0}^{K} \lambda^{p,\pi}_{n+m,n}\hat{f}^\pi_{n+m}, \qquad n \geq 0 \tag{2.10}$$

where the $\lambda^{p,\pi}_{n+m,n} = C_{n,n+m}$ are the connection coefficients (2.8).

Proof To see that $C_{n,n+m} = \lambda^{p,\pi}_{n+m,n}$, take the inner product of (2.9) with π_k against the measure ω and compare with (2.8). To see that the sum is finite as written in (2.10), take the inner product of (2.9) with f against the measure ω and use (2.4). □

We note that Proposition 2.1 is different from simple transformation of a truncated expansion as in (2.7); it is a relation that expresses the *exact* spectral information of a modified measure w in terms of the original measure ω. If the sparse result (2.9) did

not hold, then the upper limit on the sum in (2.10) would be ∞. The result allows us to commute two projection operators:

$$\mathscr{P}^p_{N-K}\mathscr{P}^\pi_N = \mathscr{P}^\pi_N\mathscr{P}^p_{N-K}. \tag{2.11}$$

A straightforward observation central to this paper is: assuming a polynomial modification, the N-coefficient connection problem of Section 2.1 can be accomplished in $\mathscr{O}(NK)$ operations by virtue of (2.10); furthermore, the initial $N-K$ coefficients correspond to an exact projection.

For a given $N > 0$, we therefore let C be the $N \times (N+K)$ matrix corresponding to the modification of spectral coefficients in (2.10). We call $\widehat{C}$ the leftmost $N \times N$ principal submatrix of C, and $\widehat{C}$ effects the polynomial connection of (2.7). In much of what follows, we assume the setup of this section, to wit:

$$\left.\begin{array}{l} I \text{ is an open interval with } \omega \text{ semi-positive on } I.\ q \in \Pi_K \text{ with } K \geq 0 \\ \text{is a polynomial semi-positive on } I. \text{ The polynomial families } \pi[\omega] \\ \text{and } p[q\omega] \text{ are accordingly defined. } C \text{ and } \widehat{C} \text{ are the } N \times (N+K) \\ \text{spectral modification and } N \times N \text{ polynomial connection matrices,} \\ \text{respectively.} \end{array}\right\} \tag{2.12}$$

We will call C the 'modification' matrix, and $\widehat{C}$ the 'connection' matrix. p is called the modified family. We occasionally refer to C as C^p_π, when such a distinction becomes necessary.

3 Convergence and conditioning

This section defines operators $\mathscr{Q}^\pm$ that are continuous versions of the modification and connection matrices C and $\hat{C}$. We first establish basic L^2 convergence properties of these operators in Section 3.2. Next Proposition 3.3 establishes one of our main results that bounds the condition number of the modification matrix $\hat{C}$ as a max-min ratio of point-evaluations of the modification polynomial q. This leads to bounds on the condition number of the operators $\mathscr{Q}^\pm$ that are given in (3.7). Lastly, the aforementioned discussions reveal the observation that the Jacobi matrix can be Cholesky-factored to produce the connection matrix. This allows us to establish Proposition 3.4: that connection matrices for non-polynomial modifications can be recovered from the Cholesky factorization of the modification function evaluated at a truncated Jacobi matrix.

3.1 Preliminaries

Vectors are denoted by boldface letters, for example $\mathbf{v}$. For any rectangular $M \times N$ matrix A, we let $\sigma_j(A)$ denote the singular values of A for $j = 1,2,\ldots,\min\{M,N\}$ that are ordered $\sigma_1 \geq \sigma_2 \geq \ldots$. The Moore-Penrose generalized inverse is denoted A^+, and any other left- or right-inverse is denoted A^-. The matrix norm $\|A\|$ is defined as the norm induced by the ℓ^2 norm on vectors. A positive matrix H is a self-adjoint

matrix satisfying $\langle \mathbf{v}, H\mathbf{v}\rangle > 0$ for all $\mathbf{v}$; semi-positivity is defined analogously. Let $\kappa_2(\cdot)$ denote the condition number associated with the ℓ^2-induced matrix norm (in this section we do not utilize the leading coefficient κ_n of the orthogonal polynomial π_n). Then $\kappa_2(A) \equiv \|A\|\|A^+\| = \sigma_1(A)/\sigma_{\min\{M,N\}}$. A result we use later is:

Lemma 3.1 *Let H be a semi-positive $N \times N$ matrix with rank R and let A be an invertible matrix of the same shape. Define $M = A^T HA$. Then:*

(i) $dim\,(ker\,(H)) = dim\,(ker\,(M))$,
(ii) $\sigma_1(M)/\sigma_1(H) \leq \sigma_1^2(A)$,
(iii) $\sigma_R(M)/\sigma_R(H) \geq \sigma_N^2(A)$.

For any $N \times N$ matrix A, we denote $A[M,M]$ the $M \times M$ leading principal submatrix of A when $M \leq N$.

The exact integral defining $\hat{f}_n^\pi$ is often approximated using a quadrature rule, e.g. a Gauss rule. The approximation to the first N coefficients using the Gauss rule can then be written $\hat{\mathbf{f}}^\pi \simeq \tilde{\mathbf{f}}^\pi = (V^\pi)^T W^\pi \mathbf{f}$, where $f_n = f\,(r_n^\pi)$ and W^π is a diagonal matrix with entries $(W^\pi)_{n,n} = w_n^\pi$. V^π is a square Vandermonde-like matrix:

$$(V^\pi)_{m,n} = \pi_m\,(r_n^\pi)\,.$$

r_n^π and w_n^π come from the N-point Gauss rule. The inverse operation, transforming $\hat{\mathbf{f}}^\pi$ to $\mathbf{f}$, is accomplished by left-multiplication with V^π. The approximation using the quadrature rule is the interpolant at the Gauss nodes:

$$\mathscr{I}_N^\pi f := \sum_{n=0}^{N-1} \tilde{f}_n^\pi \pi_n$$

Due to the exactness of the Gauss quadrature rule, $\mathscr{I}_N^\pi f = \mathscr{P}_N^\pi f$ for any $f \in \Pi_{N-1}$. For $f \notin \Pi_{N-1}$, the difference between the interpolation $\mathscr{I}_N^p f$ and the projection $\mathscr{P}_N^p f$ is the aliasing error [18] and arises due to the error in the quadrature rule.

Under assumptions (2.12) we now introduce a third operator to complement the L^2_ω-projection $\mathscr{P}^\pi$ and Gauss interpolant $\mathscr{I}^\pi$:

$$\mathscr{Q}_{N+K}^p = \mathscr{P}_N^p \mathscr{I}_{N+K}^\pi.$$

$\mathscr{Q}$ produces the degree-N L^2_w-projection of the $(N+K)$-point π-Gauss interpolant. The expansion coefficients associated with this operator are denoted $\tilde{f}_n^{p;Q}$ and are computed by $\tilde{\mathbf{f}}^{p;Q} = C(V^\pi)^T W^\pi \mathbf{f}$, and the reconstruction is given by $\mathbf{f} = V^\pi C^- \tilde{\mathbf{f}}^{p;Q}$. The right-inverse C^- that is used is computed as $(\hat{C})^{-1}\tilde{\mathbf{f}}$ and then appending K zeros. I.e.,

$$C^- = \begin{bmatrix} \widehat{C}^{-1} \\ 0_{K\times N} \end{bmatrix},$$

where $0_{M\times N}$ is an $(M \times N)$-matrix of zeros. Note that while this is a right-inverse of C, it is not the right-inverse that coincides with the Moore-Penrose pseudoinverse C^+. The inverse C^- effects an exact polynomial connection to the π basis. We define $\mathscr{Q}^-$ as the operator that corresponds to $V^\pi C^-$, and $\mathscr{Q}^+$ as the operator that corresponds to $V^\pi C^+$. For $f \in \Pi_{N-1}$, $\mathscr{Q}^- f$ is the injector that identifies Π_{N-1} as a proper subspace of Π_{N+K-1}. In contrast, $\mathscr{Q}^+ f$ produces the unique element of Π_{N+K-1} that (a) is an element of the subspace $q\Pi_{N-1} = \{qp : p \in \Pi_{N-1}\}$ and (b) $\mathscr{P}_N^p$-projects down to f.

3.2 Elementary Convergence Theory

When I is bounded, we assert that

$$\|C\| \leq \sqrt{\max_{\bar{I}} q}, \tag{3.1}$$

regardless of the size of C and provide the proof in the next subsection. An easy result is that

Lemma 3.2 *Under assumptions* (2.12) *with I bounded and $q \in C(\bar{I})$, then if $f \in L^2_\omega$, we have*

$$\|f - \mathscr{P}^p_N f\|_w \leq \sqrt{\max_{\bar{I}} q}\|f - \mathscr{P}^\pi_N f\|_\omega$$

Proof The result is almost trivial – swapping out $\mathscr{P}^w$ for $\mathscr{P}^\omega$ is the only technicality. In this proof all vectors and matrices are infinite-dimensional. Define $\hat{\mathbf{g}}^p$ to be identical to $\hat{\mathbf{f}}^p$, except for the initial N entries, which are zero; define $\hat{\mathbf{g}}^\pi$ similarly. Then

$$\|f - \mathscr{P}^p_N f\|^2_w = (\hat{\mathbf{g}}^p)^T \hat{\mathbf{g}}^p = (\hat{\mathbf{g}}^\pi)^T C^T C \hat{\mathbf{g}}^\pi \leq \left(\max_{\bar{I}} q\right) \|f - \mathscr{P}^\pi_N f\|^2_\omega.$$

□

If I is unbounded, then the result above does not provide a useful bound. A rather unsurprising result is that if the family π has a convergent projection and Gauss interpolant, then the operator $\mathscr{Q}$ is convergent with the same rate.

Proposition 3.1 *Under assumptions* (2.12)*, further assume*

$$\|f - \mathscr{P}^\pi_{N-K} f\|_\omega \leq \varepsilon_1 \qquad\qquad \|f - \mathscr{I}^\pi_N f\|_\omega \leq \varepsilon_2.$$

Then

$$\|f - \mathscr{Q}^p_N f\|_w \leq 3\sqrt{\max_{\bar{I}} q}(\varepsilon_1 + \varepsilon_2)$$

Proof Write

$$f - \mathscr{Q}^p_N f = \underbrace{(f - \mathscr{P}^\omega_N f)}_{\text{(a)}} + \underbrace{\left(\mathscr{P}^\omega_N f - \mathscr{P}^w_{N-K}\mathscr{P}^\omega_N f\right)}_{\text{(b)}} + \underbrace{\left(\mathscr{P}^w_{N-K}\mathscr{P}^\omega_N f - \mathscr{P}^w_{N-K}\mathscr{I}^\omega_N f\right)}_{\text{(c)}}$$

In the $\|\cdot\|_w$ norm, term (a) is bounded by $\varepsilon_1\sqrt{\max q}$. Use the commutation property (2.11) and Lemma 3.2 to show that term (b) is bounded by $\varepsilon_1\sqrt{\max q}$. Finally, term (c) is bounded by $\sqrt{\max q}$ times the aliasing error of the ω-operators, which is bounded by $\varepsilon_1 + \varepsilon_2$. The triangle inequality gives the result. □

3.3 Conditioning

The operators $\mathscr{I}$ and $\mathscr{Q}$ are our chief tools for practical simulations, and the goal of this section is the compare how amenable these operators are for numerical computation. In particular, we'll be concerned with computing the inverse of the operators: that is, converting expansion coefficients into point evaluations. We will see that the conditioning of each of these problems is asymptotically similar. For $(\mathscr{I}^p)^{-1}$ the required procedure is left-multiplication by V^p, and for $\mathscr{Q}^{\pm}$ it is multiplication by $V^{\pi}C^{\pm}$

We begin by analyzing the conditioning of $\mathscr{I}^{-1}$. The fact that V^{π} is 'almost' orthogonal makes the Gauss-interpolatory matrix an attractive one for analysis:

Lemma 3.3 (Gautschi, [9])

$$\kappa_2^2(V^{\pi}) = \frac{\max_n w_n^{\pi}}{\min_n w_n^{\pi}}$$

Proof Due to the accuracy of the Gauss rule (2.6), the matrix $\sqrt{W^{\pi}}V^{\pi}$ is orthogonal. Suppressing the dependence on π, this implies $V^TV = W^{-1}$, and so the singular values of V are the inverse square roots of w_n. □

Note that the same result holds for a Gauss-Radau interpolation matrix (with the Gauss quadrature weights replaced by the respective Gauss-Radau weights). A similar result for the Gauss-Lobatto Vandermonde-like matrix can be proven.

Proposition 3.2 *Assume I is bounded. Let V_{GL}^{π} be the interpolatory Vandermonde-like matrix associated with the N-point Gauss-Lobatto quadrature rule $\left\{r_{j;GL}^{\pi}, w_{j:GL}^{\pi}\right\}_{j=1}^{N}$. Define*

$$e_N = \sum_{j=1}^{N} w_{j;GL}^{\pi}\left(\pi_{N-1}(r_{j;GL}^{\pi})\right)^2,$$

and set $R = \max\{e_N, e_N^{-1}\}$. Then

$$\frac{1}{R}\left(\frac{\max_n w_{n;GL}}{\min_n w_{n;GL}}\right) \leq \kappa_2^2(V_{GL}^{\pi}) \leq R\left(\frac{\max_n w_{n;GL}}{\min_n w_{n;GL}}\right)$$

Proof Set $V = V_{GL}^{\pi}$. N-point Gauss-Lobatto quadrature rules are exact for any polynomial $\pi \in \Pi_{2N-3}$. Define the matrix E to be the $N \times N$ identity matrix, except for the (N,N) entry, which equals e_N, and let W be a diagonal matrix with the Gauss-Lobatto weights as entries. Then the accuracy of the Gauss-Lobatto rule implies that $W^{1/2}VE^{-1/2}$ is an orthogonal matrix. Therefore $V^TE^{-1}V$ has eigenvalues $w_{n;GL}^{-1}$ and so

$$\left\|V\sqrt{E_N^{-1}}\right\| \left\|\sqrt{E_N}(V)^{-1}\right\| \equiv \kappa_2(V\sqrt{E_N^{-1}}) = \sqrt{\frac{\max_n w_{n;GL}^p}{\min_n w_{n;GL}^p}}.$$

Now

$$\|V\|\|V^{-1}\| \leq \left\|V\sqrt{E_N^{-1}}\right\| \left\|\sqrt{E_N}V^{-1}\right\| \left\|\sqrt{E_N}\right\| \left\|\sqrt{E_N^{-1}}\right\|,$$

which proves the upper bound. To prove the lower bound, use Lemma 3.1 on the relation $V^T E^{-1} V = W^{-1}$ to show that $\sigma_1(W^{-1})/\sigma_1(E^{-1}) \leq \sigma_1^2(V)$ and $\sigma_N^2(V) \leq \sigma_N(W^{-1})/\sigma_N(E^{-1})$. Noting that $1/\sigma_1(E^{-1}) = \sigma_N(E)$ gives the result. □

We have determined conditioning of the Gauss Vandermonde-like matrices, but we do not need to invert them in order to transfer expansion coefficients to point-evaluations. Instead, this matrix-vector product problem has a relative condition number of $\|V^\pi\|\|\hat{\mathbf{f}}^\pi\|/\|\mathbf{f}\|$, where $\mathbf{f}$ is a vector of point-evaluations. For the particular (and not uncommon) case of taking the polynomial $p_0(r) \in \Pi_{N-1}$ to N point evaluations at the Gauss nodes, the previous formula gives the condition number

$$\kappa\left((\mathscr{I}_N^\pi)^{-1}\right) \sim \sqrt{\frac{\beta_0}{N \min_n w_n^\pi}}, \tag{3.2}$$

and this quantity is representative of the conditioning of inverting $\mathscr{I}_N^\pi$. Having established the conditioning of $\mathscr{I}^\pi$ (for both the Gauss and Gauss-Lobatto nodes), we now turn to the conditioning of $\mathscr{D}$.

Proposition 3.3 *With assumptions* (2.12)*,*

$$\kappa_2^2(V^\pi C^+) \leq \frac{\max_n q(r_n^\pi) w_n^\pi}{\min_n q(r_n^\pi) w_n^\pi} \tag{3.3}$$

$$\kappa_2^2(C) \leq \frac{\max_n q(r_n^\pi)}{\min_n q(r_n^\pi)}, \tag{3.4}$$

where r_n^π and w_n^π come from the $(N+K)$-sized quadrature rule.

Proof Let $f \in \Pi_{N-1}$ with expansion $f = \sum_n \hat{f}_n^p p_n$. Relation (2.9) implies that if we define the coefficients $\hat{\mathbf{g}}^\pi = C^T \hat{\mathbf{f}}^\mathbf{p}$ then the expansion $\sum_n \hat{g}_n^\pi \pi_n = g = qf$. Now define the size-$(N+K)$ diagonal matrices Ω and W with entries $(\Omega)_{n,n} = q(r_n^\pi)$, and $(W)_{n,n} = w_n^\pi$; and let $V = V^\pi$. Recalling that $V^{-1} = V^T W$, then

$$CV^T W \Omega^{-1} V C^T = I,$$

where I is the $N \times N$ identity matrix. I.e., if $Q = W^{1/2}\Omega^{-1/2} V C^T$, then $Q^T Q = I$, which implies that $QQ^T = H$ is of size $(N+K)\times(N+K)$, is self-adjoint, and has N unity eigenvalues and K zero eigenvalues, and is therefore semi-positive. Defining $D^{1/2} = \Omega^{1/2} W^{1/2}$ gives the relation $WVC^T CV^T W = D^{1/2} H D^{1/2}$. Lemma 3.1 implies that there are K zero eigenvalues of $WVC^T CV^T W$ and N positive eigenvalues bounded above and below by the diagonal entries in D. VC^T only has N singular values, which implies that the N positive eigenvalues of $WVC^T CV^T W$ are the squared singular values of $CV^T W$, therefore of VC^+ and this proves (3.3).

To prove (3.4) we note that $VC^T CV^{-1} = W^{-1/2}\Omega^{1/2} H \Omega^{1/2} W^{1/2}$ and therefore $C^T C$ has the same eigenvalues as $VC^T CV^{-1}$, which has the same eigenvalues as

$\Omega^{1/2}H\Omega^{1/2}$. Similar arguments as above imply that the squared singular values of C are bounded by the diagonal entries in Ω. We have shown that the squared maximum singular value of C is bounded by $\max_n q(r_n^\pi) \leq \max_I q$, which proves (3.1). □

The reasoning in the above proof that C^T effects multiplication by q allows us to identify the connection matrix $\widehat{C}$ with a fundamental quantity of the orthogonal polynomial family π: the Jacobi matrix.

Proposition 3.4 *Under assumptions* (2.12) *then the $N \times N$ leading principal submatrix of $q(J^\pi_{N+K-1})$ is positive and admits the Cholesky decomposition*

$$q(J^\pi_{N+K-1})[N,N] = \widehat{C}^T\widehat{C}, \tag{3.5}$$

where $p = p[q\omega]$.

Proof For any $N > 0$ let $f \in \Pi_{N-1}$. Then use of the $d = 0$ polynomial recurrence (2.1) implies that the map $f(r) \mapsto \mathscr{P}^\pi_N(rf(r))$ in π-spectral space is given by the linear map J^π_N. Therefore, the map $f \mapsto \mathscr{P}^\pi_N(qf)$ is the matrix $q(J^\pi_{N+K-1})[N,N] = Q$. Writing $q(r) = \sum_{k=0}^K q_k r^k$, it is clear that Q is symmetric since J^π_N is. By the arguments in the proof of Proposition 3.3, the map $f \mapsto \mathscr{P}^\pi_N(qf)$ is also given by $\widehat{C}^T\widehat{C}$; therefore Q is this matrix, and since the diagonal elements of C are all positive (see Section 4), this proves that Q symmetric has a Cholesky decomposition and therefore must be positive. □

This result is essentially a rearrangement of the results in [19]; the authors there indirectly state this result and obtain it without the use of the Christoffel modification. That $q(J^\pi_{N+K-1})[N,N]$ is strictly positive is not trivial: $q(J_{N+K-1})$ has vanishing eigenvalues if q has a root on an $(N+K-1)$-Gauss point. The result above allows us to bound the condition number of the connection matrix $\widehat{C}$.

Corollary 3.1 *With assumptions* (2.12)*,*

$$\frac{\sigma_K(q(J^\pi))}{\sigma_N(q(J^\pi))} \leq \kappa_2^2(\widehat{C}) \leq \frac{\sigma_1(q(J^\pi))}{\sigma_{N+K-1}(q(J^\pi))} \tag{3.6}$$

where $J^\pi = J^\pi_{N+K-1}$.

Proof The eigenvalues of $q(J^\pi_{N+K-1})[N,N]$ interlace with those of $q(J^\pi_{N+K-1})$ (the eigenvalue interlacing/min-max principle). Note that $\sigma_n(q(J^\pi))$ is just an ordering of $q(r_n^\pi)$ where r_n^π comes from the quadrature rule of size $(N+K-1)$. □

Remark 3.1 Although the upper bound on $\kappa_2(\widehat{C})$ seems more strict than that of $\kappa_2(C)$, we find that C is usually better conditioned. Empirical studies therefore confirm that these are not sharp bounds; however, they seem to be quite accurate when q has no zeros on $\bar{I}$.

Note that it is quite possible for q to have a root inside the interval of approximation and then r_n^π may land on such a root for some (n,N). In this case the upper bounds (3.3), (3.4), and (3.6) become infinite and are useless.

Since $\mathscr{Q}^+$ and $\mathscr{Q}^-$ have discrete representations $V^\pi C^+$ and $V^\pi C^-$, respectively, we can now obtain a characterization of the conditioning of these operators:

$$\kappa_2^2\left(\mathscr{Q}^-\right) \leq \left(\frac{\max_n q(r_n^\pi)}{\min_n q(r_n^\pi)}\right)\left(\frac{\max_n w_n^\pi}{\min_n w_n^\pi}\right) \quad \text{(size-}(N+K-1)\text{ quadrature)} \tag{3.7a}$$

$$\kappa_2^2\left(\mathscr{Q}^+\right) \leq \left(\frac{\max_n q(r_n^\pi) w_n^\pi}{\min_n q(r_n^\pi) w_n^\pi}\right) \quad \text{(size-}(N+K)\text{ quadrature)} \tag{3.7b}$$

The general message for conditioning is that when q is strictly positive on $\bar{I}$, $\kappa_2^2(C) \leq \frac{\max_I q}{\min_I q}$ when q is a polynomial, and $\widehat{C}$ obeys this same bound.

Finally, for any continuous non-polynomial q, define $q(J_N^\pi) = Sq(\Lambda)S^T$, where $S\Lambda S^T$ is the diagonalization of J^π. A limiting argument allows us to interchange the polynomial q in (3.5) with any continuous function.

Proposition 3.5 *Assume that I is bounded, ω is semi-positive, and $q \in C(\bar{I})$ is* strictly positive *on I. Then for any $N > 0$:*

$$\lim_{K\to\infty} q(J_{N+K}^{\pi[\omega]})[N,N] = \widehat{C}^T\widehat{C}, \tag{3.8}$$

where $\widehat{C} = \widehat{C}_\pi^p$ is the $N \times N$ connection matrix with $p = p[q\omega]$.

Proof For each $K \geq 0$, choose q_K as

$$q_K = \arg\min_{\phi\in\Pi_K} \|q-\phi\|_\infty.$$

By Weierstrass approximation, $\|q_K - q\|_\infty \to 0$. Let $M = \min_{\bar{I}} q > 0$. In the remainder of this proof, only consider K large enough so that $\|q_K - q\|_\infty < M$. Let $\widehat{C}_K = \widehat{C}_\pi^{p[q_K\omega]}$. Then we need only show that

1. $\widehat{C}_K \to \widehat{C}$
2. $q_K(J_{N+K}^\pi)[N,N] \sim q(J_{N+K}^\pi)[N,N]$ for large K

$q_K \to q$ in L^1_ω since the L^1 norm is bounded by the L^∞ norm under the finite measure $\mathrm{d}\omega$. All matrix norms are equivalent for a fixed N; it is convenient to choose the matrix maximum norm for property 1 and the 2-norm for property 2. The first property is straightforward since the connection coefficients are defined by (2.8) and the first N polynomials depend continuously on the first $2N$ $L^1_{q_K\omega}$ polynomial moments, which converge uniformly to the first $2N$ $L^1_{q\omega}$ polynomial moments. To show the second property, we let J_{N+K}^π have eigenvalue decomposition $J^\pi = S\Lambda S^T$, where S is orthogonal. Then

$$\|q_K(J^\pi) - q(J^\pi)\| = \|S(q_K(\Lambda) - q(\Lambda))S^T\| \leq \|q_K(\Lambda) - q(\Lambda)\| \leq \|q_K - q\|_\infty.$$

By Proposition 3.4, $q_K(J_{N+K}^\pi) = \widehat{C}_K^T\widehat{C}_K$. Taking limits on both sides we obtain the result. □

For the case of obtaining only the modified recurrence coefficient b_0 (i.e. the integral of the modified weight), this result appears as the unnumbered theorem on p. 542 of [7].

4 Modification algorithms

Having discussed the properties of the modification and connection matrices, we now describe concrete algorithms for their determination. The first method relies on polynomial factoring of q and is based on well-known methods for determining modified recurrences. The second method only requires one to be able to evaluate q on I and uses the Cholesky factorization representation of the Jacobi matrix polynomial $q(J^{\pi})$.

We emphasize that if one only requires modified recurrence coefficients, then the methods outlined in [11], [15], or more recently [2] are likely more appropriate.

4.1 Computing the transformation – successive roots

We resume the use of κ_n as the leading coefficient for polynomial π_n. Assumptions (2.12) hold as usual. We assume that the recurrence of π is known, but that of p is unknown. We further concentrate here on formation of the modification matrix C, knowing that the connection matrix $\widehat{C}$ is just a truncated version of C.

As mentioned earlier, the problem of computing information about the modified polynomials p is well-explored. However, most discussions of Christoffel transformation modifications (when q is a polynomial) center around computing the recurrence coefficients of $p[w]$. See for instance [6] and [15], with a good overview in [10]. However, in our case we are only tangentially interested in the recurrence coefficients; we instead require the modification coefficients $C_{n,n+k}$. We use Greek letters for the quantities in Section 2 associated with π (α,β,τ,κ), and Roman letters for the quantities associated with p (a,b,t,k, respectively). We also recall that any tilde'd quantities (e.g. $\tilde{\pi}_n$) refer to the same quantities, but associated with the polynomials of monic normalization.

We will need to relate the recurrence of the modified family with the modification coefficients. Assume that the α_n and β_n are known. Due to (2.9) and (2.1), one can deduce that once the constants $C_{n,n+k}$ are known, then the modified recurrence coefficients are given by

$$a_0 = \alpha_0 + \sqrt{\beta_1}\frac{C_{0,1}}{C_{0,0}}, \quad a_n = \alpha_n + \sqrt{\beta_{n+1}}\frac{C_{n,n+1}}{C_{n,n}} - \sqrt{\beta_n}\frac{C_{n-1,n}}{C_{n-1,n-1}}, \quad (n>0) \tag{4.1a}$$

$$b_0 = \beta_0 C_{0,0}^2, \qquad b_n = \beta_n \frac{C_{n,n}^2}{C_{n-1,n-1}^2}, \quad (n>0). \tag{4.1b}$$

Hence the main- and super-diagonal of C contain all the information needed to extract the modified recurrence coefficients.

Turning to the problem of computing the modification coefficients, the Christoffel theorem defining the $N\times(N+K)$ matrix entries $C_{n,n+k}$ for the degree-K polynomial q allows for computation of the coefficients explicitly via solution of N different $K\times K$ linear systems. That linear systems with known entries can be solved via Cramer's rule allows for explicit determinantal formulas to be obtained. Indeed, these explicit formulas are the statement of Christoffel's theorem. However, when K is even moderately large, it becomes computationally practical to employ a decomposition of q into products of elementary factors positive on I.

Consider the three elementary special cases:

1. $q(r) = \pm(r-z)$ for some real-valued $z \notin I$ (the sign is chosen so that q is positive on I)
2. $q(r) = (r-z)(r-z^*)$ for some $z \in \mathbb{C}\backslash\mathbb{R}$ (the asterisk * denotes complex conjugation)
3. $q(r) = (r-z)^2$ for $z \in I$

If we can compute the sparse connection matrix for each of these three cases individually, then any polynomial q semi-positive on I can be expressed as a product of the above and we can construct the full connection matrix by taking products of sparse upper-triangular matrices. All the elementary cases above are well-studied, albeit in the context of producing the modified recurrence coefficients and not the Christoffel matrix, which is more important for us.

In the first case, we have that $\pm(r-z)\tilde{p}_n = \tilde{\pi}_{n+1} \mp \tilde{\tau}_n(z)\tilde{\pi}_n$. With this relation, knowing that neither π_{n+1} nor π_n can have a zero at $r = z \notin I$, then $\mp\tilde{\tau}_{n+1}(z) = \mp\tilde{\pi}_{n+1}(z)/\tilde{\pi}_n(z)$ is positive. The orthonormal version of this modification is:

$$\pm(r-z)p_n = \pm\frac{U_n}{\sqrt{\mp\tau_{n+1}(z)}}\pi_{n+1} + U_n\sqrt{\mp\tau_{n+1}(z)}\pi_n, \quad U_n^2 = \frac{\kappa_n}{\kappa_{n+1}} = \sqrt{\beta_{n+1}} \tag{4.2}$$

For the second case, we write $(r-z)(r-z^*)\tilde{p}_n = \tilde{\pi}_{n+2} + \tilde{s}_n\tilde{\pi}_{n+1} + \tilde{r}_n\tilde{\pi}_n$. This yields a linear system for the $\tilde{r}_n$, $\tilde{s}_n$. The orthonormal version of the modification is given by

$$(r-z)(r-z^*)p_n = U_nR_n\pi_n + \frac{U_nS_n}{R_n}\pi_{n+1} + \frac{U_n}{R_n}\pi_{n+2}, \quad U_n^2 = \frac{\kappa_n}{\kappa_{n+2}} = \sqrt{\beta_{n+1}\beta_{n+2}} \tag{4.3}$$

$$R_n^2 = \frac{\mathrm{Im}\{\pi_{n+2}(z)\pi_{n+1}^*(z)\}}{\mathrm{Im}\{\pi_{n+1}(z)\pi_n^*(z)\}} = |\tau_{n+1}(z)|^2\frac{\mathrm{Im}\{\tau_{n+2}(z)\}}{\mathrm{Im}\{\tau_{n+1}(z)\}}$$
$$S_n = \frac{\mathrm{Im}\{\pi_{n+2}^*(z)\pi_n(z)\}}{\mathrm{Im}\{\pi_{n+1}(z)\pi_n^*(z)\}} = -\frac{\mathrm{Im}\{\tau_{n+1}(z)\tau_{n+2}(z)\}}{\mathrm{Im}\{\tau_{n+1}(z)\}}$$

That R_n^2 as defined is strictly positive can be deduced from the Christoffel-Darboux identity (2.2), which also ensures that the denominator of S_n never vanishes.

Finally, $(r-z)^2\tilde{p}_n = \tilde{\pi}_{n+2} + \tilde{s}_n\tilde{\pi}_{n+1} + \tilde{r}_n\tilde{\pi}_n$ with $z \in I$. Since the right-hand side must have a second-order root at $r = z$, this determines linear conditions for the unknown coefficients. The orthonormal version of the modification is given by

$$(r-z)^2p_n = U_nR_n\pi_n + \frac{U_nS_n}{R_n}\pi_{n+1} + \frac{U_n}{R_n}\pi_{n+2}, \quad U_n^2 = \frac{\kappa_n}{\kappa_{n+2}} = \sqrt{\beta_{n+1}\beta_{n+2}}. \tag{4.4}$$

$$
\begin{aligned}
R_n^2 &= \frac{\pi_{n+2}(z)\pi'_{n+1}(z) - \pi_{n+1}(z)\pi'_{n+2}(z)}{\pi_{n+1}(z)\pi'_n(z) - \pi_n(z)\pi'_{n+1}(z)} \\
&= \sqrt{\frac{\beta_{n+1}}{\beta_{n+2}}}\left[1 + \frac{\pi_{n+1}^2(z)}{\sum_{q=0}^n \pi_q^2(z)}\right], \\
S_n &= \frac{\pi_n(z)\pi'_{n+2}(z) - \pi_{n+2}(z)\pi'_n(z)}{\pi_{n+1}(z)\pi'_n(z) - \pi_n(z)\pi'_{n+1}(z)} \\
&= -\frac{1}{\sqrt{\beta_{n+2}}}\left[(z-\alpha_{n+1}) + \sqrt{\beta_{n+1}}\frac{\pi_n(z)\pi_{n+1}(z)}{\sum_{q=0}^n \pi_q^2(z)}\right],
\end{aligned}
$$

where the second equalities in the expressions for R_n^2 and S_n are obtained by using (2.2); if desired, one may use (2.1) to evaluate π'_n. Note that in this last case we have $z \in I$ so that using the ratios τ_n is not practical since we cannot guarantee that π_n does not vanish at z.

In the case where only the modified recurrence coefficients are desired, another route may be taken: there is a strong connection between Christoffel transformations and numerical linear algebra that allows one to compute the recurrence coefficients for all of the elementary cases [6], [15]. However, if we are interested in the connection coefficients, the method outlined above is more straightforward.

If we know a factorization of q into elementary factors, we can then apply (4.2), (4.3), and (4.4) in succession, alternating with an update of the recurrence constants using (4.1a) and (4.1b), and thus obtain a sequence of matrices with either two or three non-trivial diagonals. The product of these matrices is then the sought Christoffel matrix.

The positivity of $C_{n,n}$ is ensured since all the elementary expressions for R_n given above are positive, and $C_{n,n}$ is just the product of these factors. Positive entries $C_{n,n}$ correspond to the eigenvalues of $\hat{C}$ and therefore imply that the square matrix $\hat{C}$ is an invertible matrix, as expected since it is a polynomial connection matrix.

We have tested this algorithm with a polynomial modification of a Legendre weight $\omega \equiv 1$ on $I = (-1,1)$ using up to twenty repeated roots for each elementary factor. The process appears stable; for the special case $z = \pm 1$ repeated many times, the modified measure has known recurrences and Christoffel coefficients (it is a Jacobi polynomial measure) and the computed recurrence coefficients (which are derived from the computed entries in the Christoffel matrix) are within machine precision of the exact values. The algorithm has complexity $\mathcal{O}(K(N+K))$, which is as fast as the classical methods that produce only the modified recurrence coefficients.

4.2 Computing the transformation – Cholesky decomposition

We showed in Section 3, Proposition 3.4, that the matrix $\widehat{C}$ is simply the upper-triangular matrix in the Cholesky decomposition of a submatrix of $q(J^\pi)$ (the submatrix being positive-definite). Formation of C can then proceed by simply enlargening the required size of $\widehat{C}$ and then truncating appropriately. One advantage of this

method is that one does not need to determine the roots of the polynomial q; for example if one instead has the monomial representation

$$q(r) = \sum_{k=0}^{K} q_k r^k,$$

it is not necessary to find any of the roots, which can be beneficial if K is large. The simple algorithm is presented in Algorithm 1. Depending on K and N, the most expensive part is likely to be evaluation of $q(J^\pi)$, regardless of whether one evaluates directly, or via an eigendecomposition. In light of (4.1a) and (4.1b), one could consider an algorithm to obtain modified recurrence coefficients by first obtaining $\widehat{C}$ and then using the entries to find the modified recurrence. However, in comparison to simpler, well-known, $\mathcal{O}(N)$ algorithms, this method is likely not to be competitive. In some cases, it may also not be practical: the matrix $q(J^\pi)[N,N]$ can be arbitrarily close to indefinite if q has a root on $\bar{I}$. Because of this, the algorithm fails in double-precision if e.g. on $I = (-1,1)$ one tries to modify with $q = (1-r)^{40}$ since $q(J^\pi)$ and its truncated version are indefinite due to roundoff errors.

However, this method is attractive for general nonpolynomial q continuous and strictly positive as in Proposition 3.5. In this way, one can easily obtain not only a full connection matrix, but also recurrence coefficients for families with highly nontrivial weights. The difficulty in this case is the choice of K: i.e. how large must M be so that $q(J_M^\pi)[N,N]$ is an accurate representation in the limit? Naturally this is dependent on how well q can be approximated with a Weierstrass polynomial.

Algorithm 1 Computation of the connection matrix $\widehat{C}$ from the Jacobi matrix

Input: function q, Jacobi matrix J_{N+K-1}^π
Form eigendecomposition $J^\pi = S\Lambda S^T$ $(*)$
Compute $Q = Sq(\Lambda)S^T$
$Q \leftarrow Q[N,N]$, the $N \times N$ leading principal submatrix of Q
Output: $\widehat{C}$, the upper-triangular matrix in the Cholesky factorization of Q

$(*)$: One could instead form $Q = q(J^\pi)$ via direct matrix multiplication if $q \in \Pi_{K-2}$

5 Examples

In this section we illustrate the utility of connection and/or modification problems, and also of the conclusion of Proposition 3.5. Our first application is to the case of Jacobi polynomials: clearly the Jacobi weight $(1-r)^\gamma(1+r)^\delta$ is a polynomial modification of the constant (Legendre polynomial) weight when γ and δ are positive integers. Thus we view the entire family of Jacobi polynomials as a modified measure, and this allows us to show that the fast Fourier transform (FFT) can be used to compute expansions that are polynomial modifications of the Chebyshev family. We discuss how this opens up the possibility for fast algorithms for applications that use such families. The accuracy for modified families discussed in Proposition 3.1

also allows us to show spectral convergence of the operator $\mathscr{Q}_N$ to the identity for all Jacobi polynomial families.

We next quantitatively verify the convergence claims for modified families with a non-Jacobi modification of the Legendre polynomial family. Finally, we present a useful application for the Cholesky Algorithm 1: simple and accurate computation of connection coefficients for non-polynomial modifications.

5.1 Jacobi polynomials

We consider application of our results to the Jacobi orthogonal polynomial families, the particular class of orthogonal polynomials over $I = (-1,1)$ where $\omega = \omega^{(\gamma,\delta)} = (1-r)^\gamma(1+r)^\delta$ for $\gamma,\delta > -1$. We refer to e.g. [26] for a plethora of properties on this polynomial family. Jacobi polynomials form the basis for approximation in many settings.

In this section we define an analog of the $\mathscr{Q}^\pm$ operators for Jacobi polynomials. It is clear, for example, that $\omega^{(5,5)}$ is a polynomial modification for $\omega^{(0,0)}$ so that we may apply our results. We are particularly interested in the case of Jacobi polynomial families that are formed by modifying the Chebyshev family $(\gamma,\delta) = \left(-\frac{1}{2},-\frac{1}{2}\right)$. We first show that the conditioning and accuracy of the operators $\mathscr{Q}^\pm$ is comparable to a standard Gauss quadrature/interpolation method, so that little is lost by adopting this procedure. What is gained is the ability to use the fast Fourier transform to compute the Chebyshev system coefficients – since the connection matrix is sparse, this translates into a fast transform algorithm for Jacobi polynomial families. The ability to perform fast manipulations for Jacobi polynomial expansions allows one to immediately devise fast algorithms for approximation methods that utilize Jacobi polynomials.

One area of immediate application is the use of Jacobi polynomials as a basis for approximating the solution to differential equations. The generalized Jacobi polynomials/functions were introduced in [17] as weighted versions of Jacobi polynomials that extend the defintions of Jacobi polynomials to the parameter regimes $\gamma,\delta \leq -1$. For simplicity, we present the definition only for these parameter values:

$$g^{(\gamma,\delta)} = \omega^{(-1-\gamma,-1-\delta)} p^{(-1-\gamma,-1-\delta)} \qquad \gamma,\delta \leq -1$$

Because these functions naturally satisfy homogeneous boundary conditions of a certain order at $r = \pm 1$, they are prime candidates for the spectral approximation of solutions to differential equations with boundary conditions. Such a strategy was initally explored in [23] and extended to more general cases in [24]. For γ and δ negative, odd multiples of $\frac{1}{2}$, we see that $g^{(\gamma,\delta)}$ are comprised of polynomials from a family that is a polynomial modification of the Chebyshev system. In these cases the connection matrix can be used to effect a fast collocation-coefficient transform.

Jacobi polynomials also form the building blocks for function approximation on the infinite interval: the generalized Wiener rational functions are an orthonormal family of univariate complex-valued functions for expansions on the infinite real line. These functions are constructed from Jacobi polynomials. For the index $k \in \mathbb{Z}$ and

tunable parameter $s > \frac{1}{2}$, the generalized Wiener rational functions are defined as

$$\phi_k^{(s)}(x) = \frac{c_k}{(x-i)^s}\left[p_{|k|}^{(-1/2,s-3/2)}(r(x)) + \frac{2ix\,\mathrm{sgn}(k)}{x^2+1}p_{|k|-1}^{(1/2,s-1/2)}(r(x))\right], \qquad (5.1)$$

where $r(x) = \frac{1-x^2}{1+x^2}$, and c_k is a constant, and i is the imaginary unit. We use the convention that $p_{-1}^{(\gamma,\delta)} \equiv 0$. The parameter $s > \frac{1}{2}$ indicates rate of the decay of the functions as $x \to \pm\infty$. For fixed s, the family $\left\{\phi_k^{(s)}\right\}_{k\in\mathbb{Z}}$ is a complete, orthonormal basis for complex-valued L^2 functions on the real line. Thus the functions $\phi_k^{(s)}$ can be used for spectral expansions on the real line, with the advantage that one can choose the rate of decay of the functions to fit prior knowledge of the problem. We refer to [22] for details and discussion of these functions.

The form of the functions given in (5.1) shows that when s is a non-negative integer, the polynomials that define the functions are from Jacobi families that are polynomial modifications of the Chebyshev system – thus connection operations are sparse and this immediately reveals a fast transform for the Wiener rational functions.

A third application for symmetric Jacobi polynomials ($\gamma = \delta$) comes in suppression of the Gibbs phenomenon. It is known that performing a global spectral expansion on functions that are locally non-smooth produces the Gibbs phenomenon: spurious oscillations that pollute pointwise convergence globally on the region of approximation. However, it has been shown that one may suppress Gibbs oscillations and recover exponentially-accurate pointwise accuracy by transforming a Legendre (or Chebyshev) expansion into a Gegenbauer polynomial expansion (which are symmetric Jacobi polynomials) [16]. Hence one may combat the Gibbs phenomena by performing a measure modification and applying the connection matrix to transform an expansion.

We first apply the conditioning and accuracy results of Section 3 to Jacobi polynomials; this is discussed in Section 5.1.1. Section 5.1.2 validates the fact that using the FFT for Jacobi polynomial transforms can reduce computational time by orders of magnitude.

5.1.1 Convergence and conditioning

In this subsection we use our standard notation, but replace $\omega^{(\gamma,\delta)}$ and sometimes $p^{(\gamma,\delta)}$ by simply (γ,δ). In particular, we have the following notation: $\mathscr{P}^{(\gamma,\delta)}$, $\mathscr{I}^{(\gamma,\delta)}$, $L^2_{(\gamma,\delta)}$, $\|\cdot\|_{(\gamma,\delta)}$, and $V^{(\gamma,\delta)}$. Some additional function spaces used briefly here are the L^2 weighted Sobolev spaces $H^s_{(\gamma,\delta)}$. When $s \in \mathbb{N}_0$, $H^s_{(\gamma,\delta)}$ is the space of functions whose first s derivatives are in $L^2_{(\gamma,\delta)}$. Spaces for general $s \geq 0$ are then obtained by space interpolation [1].

One goal here is use of the fast Fourier transform (FFT) to perform discrete polynomial transforms. For the Chebyshev case $\gamma = \delta = -\frac{1}{2}$, this is a straightforward procedure using Gauss-type nodes [18]. The idea is that since modification/connection matrices $\widehat{C}$ are sparse when the modification is a polynomial $q \in \Pi_K$, it takes only

$\mathscr{O}(NK)$ operations to apply $\widehat{C}$. For all $\gamma,\delta \geq -\frac{1}{2}$, we define $\underline{\gamma},\underline{\delta}$ and G,D as

$$G = \left\lfloor \gamma + \frac{1}{2} \right\rfloor, \qquad D = \left\lfloor \delta + \frac{1}{2} \right\rfloor,$$
$$\underline{\gamma} = \gamma - G, \qquad \underline{\delta} = \delta - D,$$

where $\lfloor x \rfloor$ denotes the largest integer less than or equal to x. This implies that $\underline{\gamma},\underline{\delta} \in \left[-\frac{1}{2},\frac{1}{2}\right)$. Since $(1-r)^G(1+r)^D$ is a polynomial of degree $G+D$, then we propose the following implementation of the operator $\mathscr{Q}^{(\gamma,\delta)}$ for all Jacobi polynomial families with $\gamma,\delta \geq -\frac{1}{2}$:

1. Identify $q(r) = (1-r)^G(1+r)^D$ and set $\pi = \pi[\omega^{(\underline{\gamma},\underline{\delta})}]$
2. Compute the modification matrix C between families π and $p[q\omega^{(\underline{\gamma},\underline{\delta})}]$
3. Use $C\,(V^\pi)^T W^\pi$ to perform the polynomial transform to (γ,δ) coefficients.

Then $\mathscr{Q}^{(\gamma,\delta)}$ is the operator corresponding to the discrete transformation $C\,(V^\pi)^T W^\pi$ (if desired, $\widehat{C}$ may be used instead). If $\underline{\gamma} = \underline{\delta} = -\frac{1}{2}$ (which happens when 2γ and 2δ are odd), then application of $(V^\pi)^T$ can be replaced by an FFT, and application of $\widehat{C}$ is an $\mathscr{O}(N(G+D))$ operation; the result is an asymptotic $\mathscr{O}(N\log N)$ complexity for a discrete Jacobi polynomial transform. To investigate accuracy, we recall the following standard convergence results for Jacobi polynomials when $-1 < \gamma,\delta < 1$:

$$\left\| f - \mathscr{P}_N^{(\gamma,\delta)} f \right\|_{(\gamma,\delta)} \leq C N^{-s} \|f\|_{H^s_{(\gamma,\delta)}}, \tag{5.2a}$$

$$\left\| f - \mathscr{I}_N^{(\gamma,\delta)} f \right\|_{(\gamma,\delta)} \leq C N^{-s} \|f\|_{H^s_{(\gamma,\delta)}}, \tag{5.2b}$$

for N sufficiently large. See for example [1]. An immediate consequence of (5.2) and Proposition 3.1 is the spectral accuracy of the interpolation-like operator $\mathscr{Q}$.

Corollary 5.1 *Let $f \in H^s_{(\underline{\gamma},\underline{\delta})}$ for any $\gamma,\delta \geq -1/2$. Then for sufficiently large N:*

$$\left\| f - \mathscr{Q}_N^{(\gamma,\delta)} f \right\|_{(\gamma,\delta)} \leq C(N-G-D)^{-s} \|f\|_{H^s_{(\underline{\gamma},\underline{\delta})}}.$$

To investigate conditioning of operators in the Jacobi case, we note an approximate form for the Gauss weights, (15.3.10) in [26]:

$$w_n^{(\gamma,\delta)} \sim \frac{\pi}{N} w^{(\gamma+1/2,\delta+1/2)}(r_n^{(\gamma,\delta)}).$$

Coupled with the estimate from [5] that for any $\gamma,\delta \geq -\frac{1}{2}$, then $2\gamma^2/(2N+\gamma+\delta+1)^2$ is a lower bound for $1-r_N^{(\gamma,\delta)}$ and $2\delta^2/(2N+\gamma+\delta+1)^2$ is a lower bound for $1+r_1^{(\gamma,\delta)}$, this yields an approximation for estimate (3.2):

$$\kappa\left(\left(\mathscr{I}_N^{(\gamma,\delta)}\right)^{-1}\right) \sim K(\gamma,\delta)\sqrt{N}\,(2N+\gamma+\delta+1)^{\eta+1/2} \sim N^{\eta+1} \tag{5.3}$$

where $\eta = \max\{\gamma,\delta\}$. For larger γ,δ, the conditioning is algebraically more dependent on N.

In the special case of Jacobi polynomials where q is positive on the interior, useful bounds for conditioning of $\mathscr{Q}^+$ may be obtained.

Corollary 5.2 *If $|\gamma| \leq \frac{1}{2}$, $|\delta| \leq \frac{1}{2}$, and $G, D \in \mathbb{N}_0$; define $M = N + G + D$. Without loss, we assume $G \geq D$. Then the $N \times M$ Christoffel connection matrix C connecting the Jacobi polynomials of family (γ, δ) to those of family $(\gamma + G, \delta + D)$ satisfies*

$$\kappa_2(C) \leq 2^{-D} \sec^D\left(\frac{\pi}{4M+1}\right) \csc^G\left(\frac{\pi}{4M+1}\right) \sim \left(\frac{2}{\pi}N\right)^G 2^{G-D}$$

Proof We note that evaluation at the extremal nodes r_1 and r_M is sufficient to characterize estimate (3.4). A well-known bound for the extremal zeros of $P_M^{(\gamma,\delta)}$ for all γ, δ satisfying the assumptions is given by ([26]):

$$\cos\left(\frac{2M\pi}{2M+1}\right) \leq r_1 \leq \cos\left(\frac{(2M-1)\pi}{2M+1}\right),$$

with a similar bound for r_M. We therefore obtain

$$w^{(G,D)}(r_1) \geq 2^{G+D} \sin^{2D}\left(\frac{\pi}{4M+2}\right) \cos^{2G}\left(\frac{\pi}{4M+2}\right),$$

similarly for $w^{(G,D)}(r_M)$ Finally, a crude bound on the $[-1,1]$ maximum of $w^{(G,D)}$ is given by $w^{(G,D)} \leq 2^{G-D}$. Usage of (3.4) and simplification gives the result.

□

Assume $\gamma \geq \delta$. Since we must also apply $V^{(\underline{\gamma},\underline{\delta})}$ after inverting C, we introduce an additional factor of $N^{\underline{\gamma}+1/2}$ due to (5.3). Therefore, the entire operation of applying $\mathscr{Q}^+$ has a condition number

$$\kappa\left(\left(\mathscr{Q}_N^{(\gamma,\delta)}\right)^+\right) \lesssim N^G K(\underline{\gamma},\underline{\delta})\,(2N+\gamma+\delta+1)^{\underline{\gamma}+1/2} \sim N^{\gamma+1/2},$$

The operator $\mathscr{Q}^-$ has similar conditioning. Compare this to (5.3) for $\mathscr{I}^{-1}$: the operations have asymptotically similar condition numbers. We tabulate computed condition numbers for modifications of the Chebyshev weight in Table 5.1. We see as expected that the norm of V^p compared to the product of the norm of V^π and the condition number of C are similar quantities for modifications of the Chebyshev weight. However the operator $\mathscr{Q}^-$ has conditioning that is noticeably, but not egregiously, worse than the norm of V^p.

5.1.2 Jacobi FFT efficiency

We evaluate the computational efficiency for using the FFT for discrete Jacobi polynomial transforms. For the speed tests, we implement the algorithms in double-precision Fortran on a single-core workstation with a 3.06GHz processor and 2GB of RAM. Standard BLAS and FFTPACK routines were used for the matrix-vector multiplications and the FFT routines. The connection operation can be implemented as a variant of a BLAS routine. Reported timings are averaged over 5000 runs. All preprocessing steps (computation of Gauss quadrature, connection coefficient calculations, etc.) are *not* included in the timings.

	$G=D=2$			$G=D=5$		
N	$\|V^p\|$	$\|V^\pi\|\kappa_2(\hat{C})$	$\|V^\pi\|\kappa_2(C)$	$\|V^p\|$	$\|V^\pi\|\kappa_2(\hat{C})$	$\|V^\pi\|\kappa_2(C)$
50	5.33e+02	1.64e+03	3.60e+02	5.34e+04	3.27e+05	6.67e+03
100	2.87e+03	9.68e+03	2.03e+03	1.85e+06	1.96e+07	3.02e+05
200	1.58e+04	5.59e+04	1.15e+04	7.33e+07	1.01e+09	1.37e+07
400	8.84e+04	3.20e+05	6.51e+04	3.10e+09	4.88e+10	6.20e+08

	$G=0,D=2$			$G=0,D=5$		
N	$\|V^p\|$	$\|V^\pi\|\kappa_2(\hat{C})$	$\|V^\pi\|\kappa_2(C)$	$\|V^p\|$	$\|V^\pi\|\kappa_2(\hat{C})$	$\|V^\pi\|\kappa_2(C)$
50	1.01e+03	6.71e+03	1.41e+03	2.32e+05	1.31e+07	2.03e+05
100	5.59e+03	3.91e+04	8.05e+03	9.16e+06	6.96e+08	9.44e+06
200	3.12e+04	2.25e+05	4.58e+04	3.87e+08	3.40e+10	4.33e+08
400	1.76e+05	1.28e+06	2.60e+05	1.69e+10	1.60e+12	1.97e+10

Table 5.1 Approximate condition numbers for the inversion operators $\mathscr{I}^{-1}$ ($\|V^p\|$), $\mathscr{Q}^-$ ($\|V^\pi\|\kappa_2(\hat{C})$), and $\mathscr{Q}^+$ ($\|V^\pi\|\kappa_2(C)$). The polynomial family π is the Chebyshev system, and the modified family p corresponds to the Jacobi weight $w^{(-1/2+G,-1/2+D)}$.

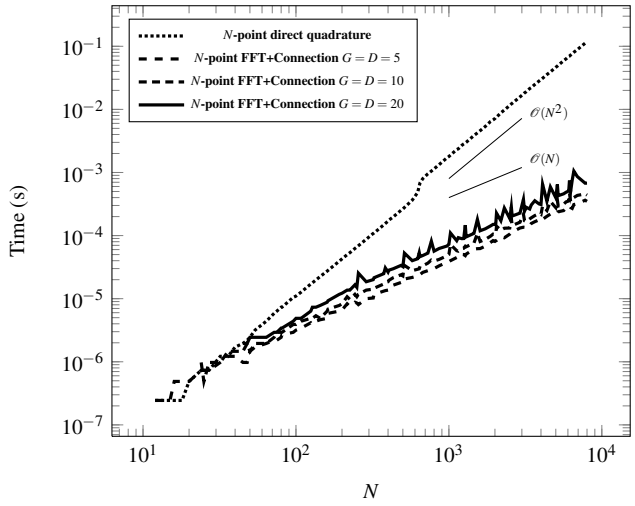

Fig. 5.1 Plot of the computational time for implementing direct (Gauss) quadrature versus usage of the FFT and the Christoffel connection matrix.

In this example, we take $\underline{\gamma} = \underline{\delta} = -1/2$ so that multiplying by V is an FFT operation. In Figure 5.1 we show the computational times required to produce expansion coefficients for $\mathscr{I}_N^{(\gamma,\delta)} f$ and $\mathscr{Q}_N^{(\gamma,\delta)} f$, where N is the number of function evaluations utilized. We illustrate the methods for $G = D$ taking values 5, 10, and 20. We see that even for very small N the time required to produce the approximation $\mathscr{Q}^{(\gamma,\delta)}$ is far less than that required to produce the Gauss interpolant $\mathscr{I}^{(\gamma,\delta)}$. In addition, the computational time to implement $\mathscr{Q}^{(\gamma,\delta)}$ is relatively insensitive to the parameters G and D.

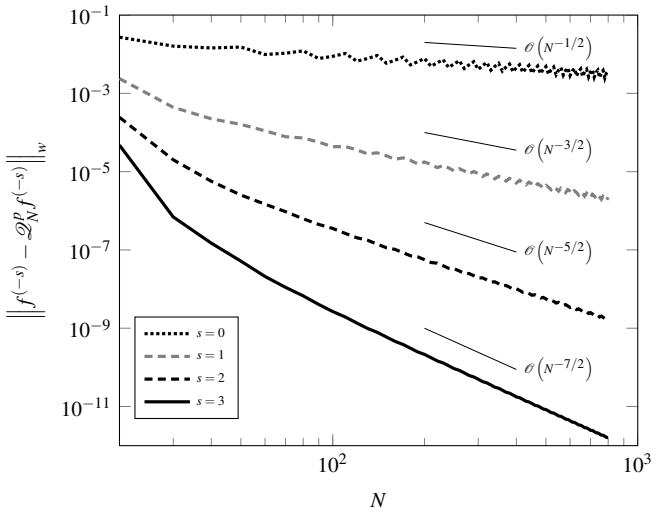

Fig. 5.2 Plot of the convergence of the operator $\mathscr{Q}_N^p$ for various values of the regularity parameter s.

5.2 Connection Accuracy

In this example we test the accuracy claims of Proposition 3.1 and Corollary 5.1. Consider the test functions on $I = (-1,1)$ defined by

$$f^{(0)} = \chi_{[-0.5,1]}, \qquad f^{(-s)} = \int_{-1}^{r} f^{(1-s)} \mathrm{d}r \qquad (s \in \mathbb{N}),$$

where χ_J is the characteristic function of the interval J. For all $s \geq 0$, $f^{(-s)} \in H_w^{s+1/2-\varepsilon}$ for all $\varepsilon > 0$ for any finite w on I. We let $\omega \equiv 1$ be the Legendre weight, and $w = q\omega$, where the modification polynomial q is given by

$$q(r) = r^2(r-0.5)^2(r+0.75)^2 \prod_{j=1}^{4} (r - z_{16} z_8^j)(r - [z_{16} z_8^j]^*) \in \Pi_{14}.$$

z_n is the n^{th} root of unity, $z_n = \exp(2\pi i/n)$, and $z_n^j = (z_n)^j$. The family $p[q\omega]$ is thus well-defined. We compute the approximations $\mathscr{Q}_N^p f^{(-s)}$ for various values of N and $s = 0,1,2,3$. The L_w^2-error between $\mathscr{Q}_N^p f^{(-s)}$ and $f^{(-s)}$ is computed using a 10^4-point Gauss-Legendre grid. According to Corollary 5.1, we expect the asymptotic convergence rate to be $N^{-1/2-s}$. This is confirmed in Figure 5.2.

5.3 Non-polynomial modifications

In many situations it is desirable to evaluate or manipulate orthogonal polynomials from a non-classical family. Such a desire comes up naturally in the field of polynomial chaos expansions in uncertainty quantification [27]. Consider a parameter X that

can be e.g. the thermal diffusion coefficient in a chemical process. For practical reasons, the value of X may not be known due to incomplete knowledge of underlying physics, or it may be tied to physically random fluctuations in a medium. In such situations it is natural to model X as a random variable with probability density function ω_X. A quantity of interest $f(X)$ (e.g. the concentration of a chemical after a reaction) is determined by physics governed by the parameter X. Evaluating f from a realization of X can involve solutions to large-scale discretized partial differential equations. It is therefore attractive to approximate the dependence of f on X with a surrogate that can be evaluated with ease. Assuming that the integral $\int f^2\omega_X \mathrm{d}x$ is finite, a natural approximation space to consider is Π_N, using elements $p_n[\omega_X]$ as basis elements: this simplifies computations and allows straightforward computation of probabilistic moments. Several extensions of this idea have been considered for vector-valued parameters [25], Karhunen-Loeve expansions [12], random fields [3], etc., and arise in applications to micro-channel fluid flow [28], electrochemical processes [4], and electromagnetic systems [20], to name a few.

A first step in applications of polynomial chaos methods is the determination of the orthogonal basis $p_n[\omega_X]$ that forms the approximation space. When ω_X corresponds to a classical orthogonal polynomial weight function, this task is accomplished with knowledge of the recurrence coefficients. However ω_X is frequently a density function generated from experimental data, or is the posterior output of a Bayesian inference problem that uses data to calibrate a posterior density. In such cases, ω_X will not be a classical weight function, and it will not be a polynomial modification of a classical measure. A simple and straightforward way to determine the recurrence coefficients for the density ω_X is to view it as a non-polynomial modification of the uniform weight function $\omega \equiv 1$, and to use Algorithm 1 to determine the ω-ω_X connection matrix, and therefore the recurrence coefficients. We therefore claim that determination of the polynomial chaos basis given an arbitrary probability density ω_X can be accomplished with little effort.

We let the unmodified measure be the Legendre measure $\omega \equiv 1$ over $I = (-1,1)$, and the modification functions be of the form

$$\omega_X(\nu, r) = \frac{1}{2} I_0(1+r) + J_0(\nu(1+r)), \tag{5.4}$$

where ν is a frequency parameter and (I_0) J_0 is the zeroth-order (modified) Bessel function of the first kind. In Figure 5.3 we plot these functions for $\nu = 10, 30, 50$. The following accuracy test is run: we pick some $N > 0$ and for various values of M we form J^{π}_{N+M}, and evaluate $\widehat{C}$ using Algorithm 1. From $\widehat{C}$ we compute the first N modified recurrence coefficients γ, δ using (4.1). These are then used with (2.1) to evaluate $p_n[\omega_X(\nu,\cdot)]$ for $0 \le n \le N-1$. Using a 10^3-point Gauss-Legendre quadrature rule, we form an approximation to the matrix with entries $A_{m,n} = \langle p_n, p_m \rangle_{\omega_X(\nu,\cdot)}$, which is the $N \times N$ identity matrix. We run the test with $N = 100$, testing orthogonality of the first 100 polynomials. For various values of M, we tabulate 2-norm errors in A in Table 5.2. This error is an indication of how accurate the computed modified recurrence coefficients are. We see that even for relatively small M and large frequency, ν, we obtain very accurate results with a very simple algorithm. This convergence occurs despite the fact that $\omega_X(J^{\pi}_{N+M})$ is not the best Weierstrass

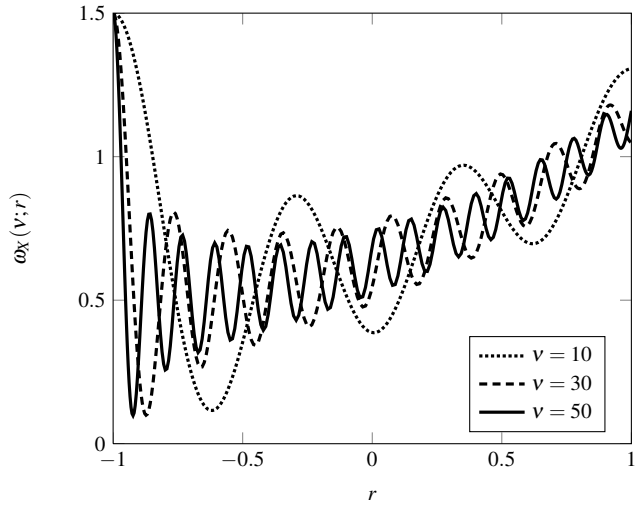

Fig. 5.3 Plot of the modification functions $\omega_X(\nu,r)$ vs. r for $\nu = 10, 30, 50$.

	$\nu = 10$	$\nu = 20$	$\nu = 30$	$\nu = 40$	$\nu = 50$
$M = 10$	2.75e-10	9.06e-04	1.21e-01	1.70e-01	2.25e-01
$M = 15$	8.18e-13	2.26e-08	1.01e-03	8.70e-02	1.44e-01
$M = 20$	8.34e-13	8.52e-13	1.24e-07	9.53e-04	6.36e-02
$M = 25$	8.43e-13	8.67e-13	9.72e-13	4.26e-07	1.50e-03
$M = 30$	8.82e-13	8.49e-13	8.26e-13	1.76e-11	1.56e-06
$M = 35$	8.52e-13	8.49e-13	8.34e-13	8.27e-13	1.81e-10
$M = 40$	7.83e-13	7.97e-13	8.09e-13	7.87e-13	7.64e-13

Table 5.2 2-norm discrete mass matrix errors for the first 100 polynomials $p_n[\omega_X(\nu,\cdot)\omega]$ generated from recurrence coefficients obtained using the Cholesky factorization algorithm with J^{π}_{100+M}.

approximant $(\omega_X)_M(J^{\pi}_{N+M})$ in the proof of Proposition 3.5 (which would be onerous to calculate). Instead, $\omega_X(J^{\pi}_{N+M})[N,N]$ is the result of using a quantity similar to $\mathscr{P}^{\pi}_M \mathscr{I}^{\pi}_{N+M} \omega_X(\nu,\cdot)$ as the approximating polynomial. In light of the convergence result (5.2b), the spectral-like accuracy in Table 5.2 is not surprising.

The results show that even for the exotic modification (5.4), Algorithm 1 is very robust, and only requires the ability to evaluate the modification function ω_X.

6 Summary

We have discussed usage of a polynomial measure modifications for producing connection coefficients, recurrence coefficients, and translating spectral expansions. The special form of polynomial modifications allows sparse connections between two polynomial families, which permits us to use well-defined spectral projections for the modified family based on information from the unmodified family. When one

of the families is the Chebyshev polynomials, we can use the FFT to first compute Chebyshev coefficients and then perform an efficient change-of-basis-like operation to produce the sought spectral expansion. This case in particular applies to Jacobi polynomials of family (γ, δ) when 2γ and 2δ are odd integers; it also applies to any family of polynomials orthogonal under the weight function $q\omega$, where q is a polynomial and ω is the Chebyshev weight.

We have investigated basic L^2 convergence theory for modifications to show that the proposed method converges with the same rate as the unmodified family approximation. For the inverse operation of obtaining point-evaluations from expansion coefficients, analysis shows that although polynomial connection problems are notoriously ill-conditioned, the fast method proposed here has conditioning that is not asymptotically worse than evaluating a traditional interpolant at the Gauss nodes using a Vandermonde-like matrix operation. Finally, we have shown that the modification polynomial evaluated at the Jacobi matrix of the unmodified family produces a positive matrix whose Cholesky decomposition furnishes the sought connection coefficients. This provides a simple method for simultaneously determining connection and modification coefficients, and the method extends in the limit to any continuous measure modification on a finite interval.

The sparsity of the connection of a polynomial measure modification can be used to effect fast transform algorithms for Jacobi polynomial spectral expansions, generalized Wiener rational functions, and the Gegenbauer reconstruction for suppression of the Gibbs phenomenon. The Cholesky decomposition for non-polynomial modifications proves useful when the measure modification is defined by functions whose explicit form it unknown. This is a common scenario in polynomial chaos expansions for uncertainty quantification when random parameter probability densities may be determined by empirical data.

Acknowledgements The authors acknowledge partial support of this work by AFOSR award FA9550-07-1-0422.

References

1. Bernardi, C., Maday, Y.: Handbook of Numerical Analysis – Techniques of Scientific Computing, vol. 5, chap. 2. Elsevier Science (1997)
2. Bueno, M.I., Dopico, F.M.: A more accurate algorithm for computing the Christoffel transformation. Journal of Computational and Applied Mathematics **205**(1), 567–582 (2007)
3. Das, S., Ghanem, R., Finette, S.: Polynomial chaos representation of spatio-temporal random fields from experimental measurements. Journal of Computational Physics **228**(23), 8726–8751 (2009)
4. Debusschere, B.J., Najm, H.N., Matta, A., Knio, O.M., Ghanem, R.G., Maître, O.P.L.: Protein labeling reactions in electrochemical microchannel flow: Numerical simulation and uncertainty propagation. Physics of Fluids **15**, 2238 (2003)
5. Erdélyi, T., Magnus, A.P., Nevai, P.: Generalized jacobi weights, christoffel functions, and jacobi polynomials. SIAM J. MATH. ANAL **25**, 602—614 (1994)
6. Galant, D.: An implemention of Christoffel's theorem in the theory of orthogonal polynomials. Mathematics of Computation **25**(113), 111–113 (1971)
7. Galant, D.: Algebraic methods for modified orthogonal polynomials. Mathematics of Computation **59**(200), 541–546 (1992)
8. Gautschi, W.: An algorithmic implementation of the generalized christoffel theorem. In: Numerical Integration, *International Series of Numerical Mathematics*, vol. 57, pp. 89–106. Birkhauser Basel (1982)

9. Gautschi, W.: The condition of vandermonde-like matrices involving orthogonal polynomials. Linear Algebra and its Applications **52-53**, 293–300 (1983)
10. Gautschi, W.: The interplay between classical analysis and (numerical) linear algebra – a tribute to gene h. golub. Electronic Transactions on Numerical Analysis **13**, 119–147 (2002)
11. Gautschi, W.: Orthogonal Polynomials: Computation and Approximation. Numerical Mathematics and Scientific Computation. Oxford University Press (2004)
12. Ghanem, R., Spanos, P.D.: Polynomial chaos in stochastic finite elements. Journal of Applied Mechanics **57**, 197 (1990)
13. Glaser, A., Liu, X., Rokhlin, V.: A fast algorithm for the calculation of the roots of special functions. SIAM Journal on Scientific Computing **29**(4), 1420–1438 (2007)
14. Golub, G., Welsch, J.: Calculation of gauss quadrature rules. Mathematics of Computation **23**(106), 221–230 (1969)
15. Golub, G.H., Kautsky, J.: Calculation of gauss quadratures with multiple free and fixed knots. Numerische Mathematik **41**(2), 147–163 (1983)
16. Gottlieb, D., Shu, C.: On the gibbs phenomenon and its resolution. SIAM Review **39**(4), 644–668 (1997)
17. Guo, B., Shen, J., Wang, L.: Generalized jacobi polynomials/functions and their applications. Applied Numerical Mathematics **59**(5), 1011–1028 (2009)
18. Hesthaven, J., Gottlieb, S., Gottlieb, D.: Spectral methods for time-dependent problems. Monographs on applied and computational mathematics. Cambridge (2007)
19. Kautsky, J., Golub, G.H.: On the calculation of jacobi matrices. Linear Algebra and its Applications **52-53**, 439–455 (1983)
20. Liu, M., Gao, Z., Hesthaven, J.S.: Adaptive sparse grid algorithms with applications to electromagnetic scattering under uncertainty. Applied Numerical Mathematics **61**(1), 24–37 (2011)
21. Maroni, P., da Rocha, Z.: Connection coefficients between orthogonal polynomials and the canonical sequence: an approach based on symbolic computation. Numerical Algorithms **47**, 291–314 (2008)
22. Narayan, A.C., Hesthaven, J.S.: A generalization of the wiener rational basis functions on infinite intervals: Part i–derivation and properties. Mathematics of Computation **80**, 1557–1557 (2011)
23. Shen, J.: Efficient Spectral-Galerkin method i. direct solvers of second- and Fourth-Order equations using legendre polynomials. SIAM Journal on Scientific Computing **15**(6), 1489 (1994)
24. Shen, J., Wang, L.: Legendre and chebyshev dual-Petrov-Galerkin methods for hyperbolic equations. Computer Methods in Applied Mechanics and Engineering **196**(37-40), 3785–3797 (2007)
25. Soize, C., Ghanem, R.: Physical systems with random uncertainties: Chaos representations with arbitrary probability measure. SIAM Journal on Scientific Computing **26**(2), 395 (2004)
26. Szegö, G.: Orthogonal Polynomials. AMS Colloquium Publications. American Mathematical Society (1939)
27. Xiu, D.: Numerical Methods for Stochastic Computations: A Spectral Method Approach. Princeton University Press (2010)
28. Xiu, D., Karniadakis, G.E.: Modeling uncertainty in flow simulations via generalized polynomial chaos. Journal of Computational Physics **187**(1), 137–167 (2003)