

Algorithm xxx: Padua2D: Lagrange Interpolation at Padua Points on Bivariate Domains

MARCO CALIARI and STEFANO DE MARCHI

University of Verona

and

MARCO VIANELLO

University of Padua

We present a stable and efficient Fortran implementation of polynomial interpolation at the “Padua points” on the square $[-1, 1]^2$. These points are unisolvent and their Lebesgue constant has minimal order of growth (log square of the degree). The algorithm is based on the representation of the Lagrange interpolation formula in a suitable orthogonal basis, and takes advantage of a new matrix formulation together with the machine-specific optimized BLAS subroutine DGEMM for the matrix-matrix product. Extension to interpolation on rectangles, triangles and ellipses is also described.

Categories and Subject Descriptors: D.3.2 [Programming Languages]: Languages Classifications—*Fortran 77*; G.1.1 [Numerical Analysis]: Interpolation; G.1.2 [Numerical Analysis]: Approximation; G.4 [Mathematics of Computing]: Mathematical Software

General Terms: Algorithms

Additional Key Words and Phrases: bivariate Lagrange interpolation, Padua points, bivariate Chebyshev orthogonal basis

1. INTRODUCTION

The problem of choosing “good” nodes on a given compact set is a central one in multivariate polynomial interpolation. Besides unisolvence, which is by no means an easy problem, for practical purposes one needs slow growth of the Lebesgue constant and computational efficiency; see, e.g., [Bojanov and Xu 2003; Carnicer et al. 2006; de Boor et al. 2000; Gasca and Sauer 2000; Sauer 1995; Xu 1996] and references therein.

In [Caliari et al. 2005] a new set of points for polynomial interpolation on the square $[-1, 1]^2$, called “Padua points”, was introduced and experimentally studied. In [Bos et al. 2006b; Bos et al. 2007] it has been proved that they are unisolvent in the space Π_n^2 of bivariate polynomials of (total) degree $\leq n$, and that they give the

Authors’ addresses: M. Caliarì and S. De Marchi, Department of Computer Science, University of Verona, Strada Le Grazie 15, 37134 Verona (Italy); email: marco.caliari@univr.it and stefano.demarchi@univr.it; M. Vianello, Department of Pure and Applied Mathematics, University of Padua, Via Trieste 63, 35121 Padova (Italy); email: marcov@math.unipd.it.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 20YY ACM 0098-3500/20YY/1200-0001 \$5.00

first known example of non tensor-product optimal interpolation in two variables, since their Lebesgue constant has minimal order of growth of $\mathcal{O}((\log n)^2)$. Moreover, an explicit representation of their fundamental Lagrange polynomials was given, which is based on the reproducing kernel of the space $\Pi_n^2([-1, 1]^2)$ equipped with the bivariate Chebyshev inner product (cf. [Dunkl and Xu 2001; Reimer 2003]).

In this paper we present a stable and efficient Fortran implementation of the Lagrange interpolation formula at the Padua points, with cost $\mathcal{O}(n^3)$ flops for the evaluation (once and for all) of the coefficients of the interpolation polynomial represented in the Chebyshev orthonormal basis, plus an additional cost of $\mathcal{O}(n^2)$ flops for the evaluation at each target point. Remarkable speedups with respect to the existing algorithm used in [Caliari et al. 2007] are obtained by a new matrix formulation. The effect of machine-specific optimized versus standard BLAS libraries is also investigated.

2. PADUA POINTS

There are four families of Padua points (cf. [Caliari et al. 2007]), which correspond to successive 90 degrees rotations (clockwise for degree n even and counterclockwise for n odd) of the square $[-1, 1]^2$. Here we describe one of them (also termed the “first family” of Padua points) as well as the associated Lagrange interpolation formula. Interpolation for the other families can be easily obtained by composition with the corresponding rotation, as described in Sec. 3.4.

The $N := (n + 1)(n + 2)/2 = \dim(\Pi_n^2)$ Padua points corresponding to degree n are the set of points

$$\text{Pad}_n := \{\boldsymbol{\xi} = (\xi_1, \xi_2)\} := \left\{ \gamma \left(\frac{k\pi}{n(n+1)} \right), \quad k = 0, \dots, n(n+1) \right\}, \quad (1)$$

where $\gamma(t)$ is their “generating curve” (cf. [Bos et al. 2006b])

$$\gamma(t) := (-\cos((n+1)t), -\cos(nt)), \quad t \in [0, \pi]. \quad (2)$$

Notice that two of the points are consecutive vertices of the square, $2n - 1$ other points lie on the edges of the square, and the remaining (interior) points are double points corresponding to self-intersections of the generating curve (see Fig. 1). Clearly, such structure remains invariant for the other three families up to the corresponding rotation.

Denote by C_{n+1} the set of the $n + 1$ Chebyshev–Gauss–Lobatto points

$$C_{n+1} := \{z_j^n := \cos(j\pi/n), \quad j = 0, \dots, n\}. \quad (3)$$

The Padua points (for n even) were first introduced in [Caliari et al. 2005, formula (9)] (in that formula there is a misprint, $n - 1$ has to be replaced by $n + 1$). For the sake of clarity, we give again that definition:

$$\text{Pad}_n := \{\boldsymbol{\xi} = (\mu_j, \eta_k), \quad 0 \leq j \leq n; \quad 0 \leq k \leq n/2\}$$

where

$$\mu_j = z_j^n, \quad \eta_k = \begin{cases} z_{2k}^{n+1} & j \text{ odd} \\ z_{2k+1}^{n+1} & j \text{ even} \end{cases}$$

It is clear that $\mu_j \in C_{n+1}$, $\eta_k \in C_{n+2}$, and denoting by C_{n+1}^{even} , C_{n+1}^{odd} the restrictions of C_{n+1} to even and odd indexes, then

$$\text{Pad}_n = (C_{n+1}^{\text{odd}} \times C_{n+2}^{\text{even}}) \cup (C_{n+1}^{\text{even}} \times C_{n+2}^{\text{odd}}) \subset C_{n+1} \times C_{n+2},$$

which is valid also for n odd.




Fig. 1. The Padua points with their generating curve for $n = 12$ (left, 91 points) and $n = 13$ (right, 105 points), also as union of two Chebyshev-like grids (open bullets and filled bullets).

The fundamental Lagrange polynomials of the Padua points are

$$L_{\xi}(\mathbf{x}) = w_{\xi}(K_n(\xi, \mathbf{x}) - T_n(\xi_1)T_n(x_1)), \quad L_{\xi}(\eta) = \delta_{\xi\eta}, \quad \xi, \eta \in \text{Pad}_n, \quad (4)$$

where $K_n(\mathbf{x}, \mathbf{y})$, $\mathbf{x} = (x_1, x_2)$ and $\mathbf{y} = (y_1, y_2)$, is the reproducing kernel of the space $\Pi_n^2([-1, 1]^2)$ equipped with the inner product

$$\langle f, g \rangle := \frac{1}{\pi^2} \int_{[-1, 1]^2} f(x_1, x_2)g(x_1, x_2) \frac{dx_1}{\sqrt{1-x_1^2}} \frac{dx_2}{\sqrt{1-x_2^2}}, \quad (5)$$

that is

$$K_n(\mathbf{x}, \mathbf{y}) = \sum_{k=0}^n \sum_{j=0}^k \hat{T}_j(x_1)\hat{T}_{k-j}(x_2)\hat{T}_j(y_1)\hat{T}_{k-j}(y_2). \quad (6)$$

Here \hat{T}_j denotes the normalized Chebyshev polynomial of degree j , i.e. $\hat{T}_0 = T_0 \equiv 1$, $\hat{T}_p = \sqrt{2}T_p$, $T_p = \cos(p \arccos(\cdot))$. Moreover, the weights w_{ξ} are

$$w_{\xi} := \frac{1}{n(n+1)} \cdot \begin{cases} 1/2 & \text{if } \xi \text{ is a vertex point} \\ 1 & \text{if } \xi \text{ is an edge point} \\ 2 & \text{if } \xi \text{ is an interior point} \end{cases} \quad (7)$$

We notice that the $\{w_{\xi}\}$ are indeed weights of a cubature formula for the product Chebyshev measure, which is exact on “almost all” polynomials in $\Pi_{2n}^2([-1, 1]^2)$, namely on all polynomials orthogonal to $T_{2n}(x_1)$. Such a cubature formula stems

from quadrature along the generating curve and is the key to obtaining the fundamental Lagrange polynomials (4); cf. [Bos et al. 2006b]. A more abstract approach to the construction of the Lagrange polynomials, based on the theory of polynomial ideals and multivariate orthogonal polynomials, can be found in [Bos et al. 2007].

The polynomial interpolation formula can be written, in view of (4) and (6), in the bivariate Chebyshev orthonormal basis as

$$\begin{aligned} \mathcal{L}_n f(\mathbf{x}) &= \sum_{\boldsymbol{\xi} \in \text{Pad}_n} f(\boldsymbol{\xi}) w_{\boldsymbol{\xi}} (K_n(\boldsymbol{\xi}, \mathbf{x}) - T_n(\xi_1) T_n(x_1)) \\ &= \sum_{k=0}^n \sum_{j=0}^k c_{j,k-j} \hat{T}_j(x_1) \hat{T}_{k-j}(x_2) - \sum_{\boldsymbol{\xi} \in \text{Pad}_n} f(\boldsymbol{\xi}) w_{\boldsymbol{\xi}} \hat{T}_n(\xi_1) \frac{\hat{T}_0(\xi_2)}{\sqrt{2}} \hat{T}_n(x_1) \frac{\hat{T}_0(x_2)}{\sqrt{2}} \\ &= \sum_{k=0}^n \sum_{j=0}^k c_{j,k-j} \hat{T}_j(x_1) \hat{T}_{k-j}(x_2) - \frac{c_{n,0}}{2} \hat{T}_n(x_1) \hat{T}_0(x_2), \end{aligned} \tag{8}$$

where the coefficients

$$c_{j,k-j} := \sum_{\boldsymbol{\xi} \in \text{Pad}_n} f(\boldsymbol{\xi}) w_{\boldsymbol{\xi}} \hat{T}_j(\xi_1) \hat{T}_{k-j}(\xi_2), \quad 0 \leq j \leq k \leq n \tag{9}$$

can be computed once and for all.

3. THE INTERPOLATION ALGORITHM

The following definition will be useful below. Given a vector $S = (s_1, \dots, s_m) \in [-1, 1]^m$, define the rectangular Chebyshev matrix

$$\mathbb{T}(S) := \begin{pmatrix} \hat{T}_0(s_1) & \cdots & \hat{T}_0(s_m) \\ \vdots & \cdots & \vdots \\ \hat{T}_n(s_1) & \cdots & \hat{T}_n(s_m) \end{pmatrix} \in \mathbb{R}^{(n+1) \times m}. \tag{10}$$

3.1 Old algorithm

In the recent paper [Caliari et al. 2007] the interpolation formula (8)–(9) has been implemented by the following factorization of the coefficient matrix, which was already used for hyperinterpolation at the Morrow–Patterson–Xu points in [Caliari et al. 2006]. Considering the vectors of the one-dimensional projections of the Padua points and the corresponding Chebyshev matrices

$$\mathbb{T}_i := \mathbb{T}(\text{Pad}_n^i) \in \mathbb{R}^{(n+1) \times N}, \quad \text{Pad}_n^i := (\xi_i)_{\boldsymbol{\xi} \in \text{Pad}_n}, \quad i = 1, 2, \tag{11}$$

and the $N \times N$ diagonal matrix $\mathbb{D}(f)$, having the vector $(w_{\boldsymbol{\xi}} f(\boldsymbol{\xi}))_{\boldsymbol{\xi} \in \text{Pad}_n}$ in the main diagonal, following [Caliari et al. 2006] it is easy to prove that the coefficients $c_{j,k-j}$ in (9) are the upper-left triangular part of the matrix

$$(c_{p,q}) = \mathbb{C}(f) := \mathbb{T}_1 \mathbb{D}(f) \mathbb{T}_2^t \in \mathbb{R}^{(n+1) \times (n+1)}. \tag{12}$$

Indeed, we have that $c_{p,q}$ in (12) is exactly $c_{j,k-j}$ in (9) with $j = p$, $k = p + q$, $0 \leq j \leq k \leq n$. Notice that the multiplication by the diagonal matrix $\mathbb{D}(f)$ is simply a scaling of the columns of \mathbb{T}_1 or of the rows of \mathbb{T}_2 . Then, defining the

Fourier–Chebyshev coefficient matrix

$$\mathbb{C}_0(f) := \begin{pmatrix} c_{0,0} & c_{0,1} & \cdots & \cdots & c_{0,n} \\ c_{1,0} & c_{1,1} & \cdots & c_{1,n-1} & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ c_{n-1,0} & c_{n-1,1} & 0 & \cdots & 0 \\ c_{n,0}/2 & 0 & \cdots & 0 & 0 \end{pmatrix} \in \mathbb{R}^{(n+1) \times (n+1)}, \quad (13)$$

it is easy to see that the polynomial interpolation formula (8) can be evaluated at any $\mathbf{x} = (x_1, x_2) \in [-1, 1]^2$ by

$$\mathcal{L}_n f(\mathbf{x}) = (\mathbb{T}(x_1))^t \mathbb{C}_0(f) \mathbb{T}(x_2), \quad (14)$$

cf. (10), taking suitably into account the triangular structure of $\mathbb{C}_0(f)$ (notice that $\mathbb{T}(x_i)$, $i = 1, 2$, is a column vector).

3.2 New algorithm

A substantial improvement can be obtained by a different matrix factorization. With a little abuse of notation, we also denote by $C_{n+1} = (z_0^n, \dots, z_n^n)$ the ordered vector of the Chebyshev–Gauss–Lobatto points. Consider the matrices (cf. (10))

$$\mathbb{P}_1 := \mathbb{T}(C_{n+1}) \in \mathbb{R}^{(n+1) \times (n+1)}, \quad \mathbb{P}_2 := \mathbb{T}(C_{n+2}) \in \mathbb{R}^{(n+1) \times (n+2)}, \quad (15)$$

and define the $(n+1) \times (n+2)$ matrix computed corresponding to the Chebyshev-like grid $C_{n+1} \times C_{n+2}$

$$\mathbb{G}(f) = (g_{r,s}) := \begin{cases} w_{\boldsymbol{\eta}} f(z_r^n, z_s^{n+1}) & \text{if } \boldsymbol{\eta} = (z_r^n, z_s^{n+1}) \in \text{Pad}_n \\ 0 & \text{if } \boldsymbol{\eta} = (z_r^n, z_s^{n+1}) \in (C_{n+1} \times C_{n+2}) \setminus \text{Pad}_n \end{cases} \quad (16)$$

Then, it is easy to check that the matrix $\mathbb{C}(f)$ in (12) can also be computed by the alternative factorization

$$\mathbb{C}(f) = \mathbb{P}_1 \mathbb{G}(f) \mathbb{P}_2^t. \quad (17)$$

The evaluation of the polynomial interpolant is now performed as in (14).

The computational cost of the coefficient matrix using (17) is $\mathcal{O}(n^3)$, whereas that of (12) is $\mathcal{O}(n^4)$. Such an improvement is due to the fact that the relevant matrices have $\mathcal{O}(n^2)$ instead of $\mathcal{O}(n^3)$ elements. The matrix formulation (14)–(17) of the interpolation problem will be the computational key for the Fortran implementation described in Sec. 4.

3.3 Error estimates

In view of the growth estimate for the Lebesgue constant proved in [Bos et al. 2006b] and of the multivariate extension of Jackson’s theorem (cf., e.g., [Bagby et al. 2002] and references therein), we have that for $f \in C^p([-1, 1]^2)$, $p > 0$,

$$\|f - \mathcal{L}_n f\|_{\infty} \leq (1 + \Lambda_n) E_n(f) \leq \mathcal{C}(f; p) (1 + \Lambda_n) n^{-p} = \mathcal{O}(n^{-p} (\log n)^2),$$

where \mathcal{C} is a suitable constant independent of n but depending on f and p . However, this a priori estimate is essentially qualitative.

A desirable feature is the availability of a reliable and, if possible, a posteriori quantitative estimate of the error. To this purpose, we have used the fact that

interpolation at the Padua points is “close” to a discretized truncated Fourier–Chebyshev expansion. We then obtain the following a posteriori estimate of the interpolation error (cf. (8)–(9))

$$\|f - \mathcal{L}_n f\|_\infty \approx \|f - \mathcal{S}_n f\|_\infty \approx 2 \sum_{k=n-2}^n \sum_{j=0}^k |c_{j,k-j}|, \quad (18)$$

where $\mathcal{S}_n f$ is the orthogonal projection of f on $\Pi_n^2([-1, 1]^2)$ with respect to the inner product (5).

The first approximation can be rigorously justified (cf. [Caliari et al. 2007]), whereas the second is somewhat “empirical”, but reminiscent of popular error estimates for one-dimensional Chebyshev expansions. In practice, indeed, as in the one-dimensional case it happens that (18) tends to be an overestimate for smooth functions, and an underestimate for functions of low regularity (due to the fast/slow decay of the Fourier–Chebyshev coefficients). Application to classical test sets for interpolation has shown that the behavior of this error estimate is satisfactory (see [Caliari et al. 2007]), and thus it is provided as an output of our interpolation subroutine.

In order to give a flavour of the numerical performance of the interpolation algorithm and reliability of the error estimate (18), we report in Table I some results on the classical Franke test suite (cf. [Franke 1982]) taken from [Caliari et al. 2007]. We recall that all errors are normalized to the maximum deviation of the function from its mean, and that the “true” errors have been computed on a 100×100 uniform control grid.

degree n	F_1	F_2	F_3	F_4	F_5	F_6
20	2E-2 (2E-2)	6E-2 (8E-2)	1E-5 (8E-5)	7E-10 (1E-7)	6E-5 (8E-4)	4E-8 (5E-7)
40	2E-6 (1E-5)	2E-3 (2E-3)	2E-11 (2E-10)	8E-15 (1E-15)	4E-13 (2E-11)	1E-14 (1E-13)
60	2E-11 (3E-10)	6E-5 (7E-5)	3E-14 (5E-15)	2E-14 (5E-15)	8E-15 (1E-15)	3E-14 (6E-15)

Table I. “True” and estimated (in parentheses) normalized errors of interpolation of Padua points for the classical Franke test suite.

3.4 Mapping the square

The interpolation formula (14) can be extended to: a) the other three families of Padua points and b) other domains, by means of a suitable mapping of the square. Indeed, when a smooth and surjective transformation

$$\sigma: [-1, 1]^2 \rightarrow \Omega, \quad \mathbf{t} \mapsto \mathbf{x} = \sigma(\mathbf{t}), \quad \Omega \subset \mathbb{R}^2 \quad (19)$$

is given, we can construct the (in general nonpolynomial) interpolation formula on Ω ,

$$\mathcal{L}_n f(\mathbf{x}) = \mathcal{L}_n f(\mathbf{x}; \sigma, \Omega) := \mathbb{T}(\sigma_1^{-1}(\mathbf{x}))^\dagger \mathbb{C}_0(f \circ \sigma) \mathbb{T}(\sigma_2^{-1}(\mathbf{x})). \quad (20)$$

Here we denote by $\sigma_i^{-1}(\mathbf{x})$, $i = 1, 2$ the components of the “inverse” transformation, which in general has to be defined by choosing a point in the inverse image of

\mathbf{x} (recall that σ is assumed to be surjective but not necessarily injective). Any particular choice will give a valid interpolation formula.

For example, polynomial interpolation over any nondegenerate parallelogram can be obtained by a suitable invertible affine map such as

$$\sigma(\mathbf{t}) := A\mathbf{t} + \mathbf{v}, \quad \sigma^{-1}(\mathbf{x}) = A^{-1}(\mathbf{x} - \mathbf{v}), \quad A \in \mathbb{R}^{2 \times 2}. \quad (21)$$

As a special case, interpolation over the reference square rotated by an angle θ can be immediately obtained from

$$A := \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix}, \quad A^{-1} = A^t, \quad \mathbf{v} = \mathbf{0}. \quad (22)$$

In particular, Lagrange interpolation at the other three families of Padua points (say, the “second”, “third” and “fourth” family), corresponds to $\theta = -\pi/2, -\pi, -3\pi/2$ for even degree n , and $\theta = \pi/2, \pi, 3\pi/2$ for odd degree, respectively (here $\Omega = [-1, 1]^2$). On the other hand, interpolation of a function f defined on a rectangle such as $\Omega = R(\mathbf{a}, \mathbf{b}) := [a_1, b_1] \times [a_2, b_2]$, corresponds to

$$A := \frac{1}{2} \begin{pmatrix} b_1 - a_1 & 0 \\ 0 & b_2 - a_2 \end{pmatrix}, \quad \mathbf{v} = \frac{1}{2}(\mathbf{b} + \mathbf{a}). \quad (23)$$

The examples above concern transformations that are globally invertible and still generate a polynomial interpolant. But it is also possible to use transformations into bivariate domains with different geometric structures, provided that they are surjective and smooth in order to preserve as much as possible the polynomial approximation quality. These are the cases, for example, of triangles, generalized rectangles and generalized sectors, where we finally get a nonpolynomial interpolant.

As for the case of a nondegenerate triangle $\Omega = T(\mathbf{u}, \mathbf{v}, \mathbf{w})$ with vertices $\mathbf{u} = (u_1, u_2)$, $\mathbf{v} = (v_1, v_2)$ and $\mathbf{w} = (w_1, w_2)$, it is convenient to use the Proriot (also known as the Duffy) quadratic map

$$\sigma(\mathbf{t}) := \frac{1}{4}(\mathbf{v} - \mathbf{u})(1 + t_1)(1 - t_2) + \frac{1}{2}(\mathbf{w} - \mathbf{u})(1 + t_2) + \mathbf{u}. \quad (24)$$

Notice that this is non injective, since it maps the whole upper side of the square ($t_2 = 1$) to the vertex \mathbf{w} . An “inverse” can be defined via the auxiliary map $\rho(\mathbf{x})$ from $T(\mathbf{u}, \mathbf{v}, \mathbf{w})$ to the unit triangle with vertices $(0, 0)$, $(1, 0)$ and $(0, 1)$

$$\begin{aligned} \rho_1(\mathbf{x}) &:= \frac{x_1(w_2 - u_2) - x_2(w_1 - u_1) + (w_1 - u_1)u_2 - (w_2 - u_2)u_1}{(v_1 - u_1)(w_2 - u_2) - (v_2 - u_2)(w_1 - u_1)} \\ \rho_2(\mathbf{x}) &:= \frac{x_1(v_2 - u_2) - x_2(v_1 - u_1) + (v_1 - u_1)u_2 - (v_2 - u_2)u_1}{(w_1 - u_1)(v_2 - u_2) - (w_2 - u_2)(v_1 - u_1)}, \end{aligned}$$

and by choosing a point in the inverse image of \mathbf{w} , e.g. as

$$\begin{aligned} \sigma^{-1} &: T(\mathbf{u}, \mathbf{v}, \mathbf{w}) \rightarrow [-1, 1]^2 \\ \sigma^{-1}(\mathbf{x}) &:= \begin{cases} (0, 1) & \mathbf{x} = \mathbf{w} \ (\rho_2 = 1) \\ \left(\frac{2\rho_1}{1 - \rho_2} - 1, 2\rho_2 - 1 \right) & \mathbf{x} \neq \mathbf{w} \end{cases} \quad (25) \end{aligned}$$

The resulting interpolant (20) at the Padua points mapped into the triangle becomes rational.

As an example of the use of polar-like coordinates, consider an ellipse $\Omega = E(\mathbf{c}, \alpha, \beta)$ centered in $\mathbf{c} = (c_1, c_2)$ and with x_1 -semiaxis α and x_2 -semiaxis β . It is convenient (cf. [Bos et al. 2006a]) to define the starlike-polar map

$$\sigma_1(\mathbf{t}) := c_1 - \alpha t_2 \sin\left(\frac{\pi}{2} t_1\right), \quad \sigma_2(\mathbf{t}) := c_2 + \beta t_2 \cos\left(\frac{\pi}{2} t_1\right), \quad (26)$$

which distributes the interpolation points in a more symmetric way with respect to standard polar coordinates (that cluster points near the right horizontal semiaxis). The map is not globally invertible, since the whole segment $t_2 = 0$ is mapped to the ellipse center \mathbf{c} , and $\sigma(1, t_2) = \sigma(-1, -t_2) = (c_1 - \alpha t_2, c_2)$ for every $t_2 \in [-1, 1]$. An inverse can be defined via the angle

$$\theta(\mathbf{x}) := \arctan\left(\frac{\beta c_1 - x_1}{\alpha x_2 - c_2}\right), \quad \mathbf{x} \neq \mathbf{c}, \quad (27)$$

and by choosing a point in the relevant inverse images, e.g. as

$$\sigma^{-1}: E(\mathbf{c}, \alpha, \beta) \rightarrow [-1, 1]^2$$

$$\sigma^{-1}(\mathbf{x}) := \begin{cases} \left(1, \frac{c_1 - x_1}{\alpha}\right) & x_2 = c_2 \\ \left(\frac{2}{\pi} \theta, \frac{x_2 - c_2}{\beta \cos \theta}\right) & x_2 \neq c_2 \end{cases} \quad (28)$$

From the identity $1/\cos^2 \theta = 1 + \tan^2 \theta$, we can rewrite the last row of (28) in the numerically more stable form

$$\sigma_1^{-1}(\mathbf{x}) = \frac{2}{\pi} \arctan\left(\frac{\beta c_1 - x_1}{\alpha x_2 - c_2}\right), \quad (29)$$

$$\sigma_2^{-1}(\mathbf{x}) = \text{sign}(x_2 - c_2) \frac{\sqrt{\beta^2(x_1 - c_1)^2 + \alpha^2(x_2 - c_2)^2}}{\alpha\beta}.$$

4. CODE

The software, which implements interpolation at the Padua points by the algorithm described in Sec. 3.2, is written in ANSI Fortran 77 and uses double precision. It depends on the BLAS routines `DDOT`, `DGEMM` (for the matrix-matrix product (17)), `DSCAL`, `LSAME` and `XERBLA`. Interpolation at degree n , i.e. by $N = (n+1)(n+2)/2$ Padua points, has the following storage requirements: four N -dimensional arrays, `PD1`, `PD2`, `FPD` and `WPD` containing the Padua points, the values of the function at the Padua points and the weights respectively; two auxiliary $(n+2) \times (n+2)$ arrays `RAUX1` and `RAUX2` to compute the matrix products in (17); an $(n+1) \times (n+1)$ array `C0` for the coefficient matrix (13), for an overall allocation of about $5n^2$ floating point numbers. The user-callable subprograms are the following:

PDPTS Subroutine which generates the Padua points for a given degree n and the corresponding weights, see (1)–(7).

CHEB Subroutine which computes a Chebyshev (column) vector $\mathbb{T}(s)$, $s \in [-1, 1]$, see (10), by the three-term recurrence of the Chebyshev polynomials. It is called by `PADUA2`, see (15)–(17), and by `PD2VAL`, see (14).

PADUA2 Subroutine which computes the coefficient matrix $C_0(f)$, see Sec. 3.2.

PD2VAL Function which returns the value of the interpolated function $\mathcal{L}_n f(\mathbf{x})$ at an arbitrary target point \mathbf{x} , see (14).

The software also allows for interpolation on rectangles, ellipses and triangles (see Sec. 3.4). This is described in the three example drivers `DRVREC`, `DRVTRI`, `DRVELL` that compute the interpolation of the Franke test function on the unit square $[0, 1]^2$, the unit triangle $T((0, 0), (1, 0), (0, 1))$, and the disk with center $(0.5, 0.5)$ and radius 0.5, respectively. The driver `DRVREC` also shows how to interpolate at the three other families of Padua points by composition with the corresponding rotations.

4.1 Benchmarks

In this section we compare the performances of the old and new algorithms (see Sec. 3) in the computation of the coefficient matrix C_0 , both implemented using the standard BLAS and the machine-specific optimized BLAS (ACML [AMD 2006] in our numerical experiments). For the relevant (double precision, general) matrix-matrix products, the BLAS subroutine `DGEMM` is used. In all cases we observe, from Table II that, as expected, the computation time for the evaluation of the coefficients scales as n^4 for the old algorithm (12) and as n^3 for the new algorithm (17). The speedup obtained by the new algorithm in this range of degrees (order of hundreds, i.e., several thousands interpolation points) turns out to be of 2–3 orders of magnitude. We also note that the use of optimized BLAS becomes particularly important when dealing with very high interpolation degrees.

Algorithm	Library	$n = 100$	$n = 200$	$n = 300$	$n = 400$
Old	BLAS	0.39E0	0.45E1	0.21E2	0.68E2
	ACML	0.77E-1	0.63E0	0.27E1	0.79E1
New	BLAS	0.60E-2	0.48E-1	0.17E0	0.52E0
	ACML	0.20E-2	0.10E-1	0.35E-1	0.80E-1

Table II. CPU times (seconds) for the computation of the coefficient matrix $C_0(f)$ in (13) by the old and the new algorithms in Sec. 3.

The detailed environment of our numerical experiments is as follows:

CPU AMD Opteron 248 2.2GHz, 1MB L2 Cache
 OS GNU/Linux 2.6.5
 Compiler GNU Fortran (GCC) 3.3.3
 Compiler options `-O3 -m64`
 BLAS BLAS (Basic Linear Algebra Subprograms) (cf. [Blackford et al. 2002])
 ACML AMD Core Math Library (ACML) 3.6.0 (cf. [AMD 2006])

REFERENCES

- AMD 2006. AMD Core Math Library (ACML). Version 3.6.0. Available at <http://developer.amd.com/acml.aspx>.
 BAGBY, T., BOS, L., AND LEVENBERG, N. 2002. Multivariate simultaneous approximation. *Constr. Approx.* 18, 569–577.

- BLACKFORD, L. S., DEMMEL, J., DONGARRA, J., DUFF, I., HAMMARLING, S., HENRY, G., HEROUX, M., KAUFMAN, L., LUMSDAINE, A., PETITET, A., POZO, R., REMINGTON, K., AND WHALEY, R. C. 2002. An Updated Set of Basic Linear Algebra Subprograms (BLAS). *ACM Trans. Math. Soft.* 28, 2, 135–151. Available at <http://www.netlib.org/blas/>.
- BOJANOV, B. AND XU, Y. 2003. On polynomial interpolation of two variables. *J. Approx. Theory* 120, 267–282.
- BOS, L., CALIARI, M., DE MARCHI, S., AND VIANELLO, M. 2006a. Bivariate interpolation at Xu points: results, extensions and applications. *Electron. Trans. Numer. Anal.* 25, 1–16.
- BOS, L., CALIARI, M., DE MARCHI, S., VIANELLO, M., AND XU, Y. 2006b. Bivariate Lagrange interpolation at the Padua points: the generating curve approach. *J. Approx. Theory* 143, 15–25.
- BOS, L., DE MARCHI, S., VIANELLO, M., AND XU, Y. 2007. Bivariate Lagrange interpolation at the Padua points: the ideal theory approach. *Numer. Math.* 108, 1, 43–57.
- CALIARI, M., DE MARCHI, S., AND VIANELLO, M. 2005. Bivariate polynomial interpolation on the square at new nodal sets. *Appl. Math. Comput.* 165, 2, 261–274.
- CALIARI, M., DE MARCHI, S., AND VIANELLO, M. 2007. Bivariate Lagrange interpolation at the Padua points: computational aspects. *J. Comput. Appl. Math.*, available online 23 October 2007.
- CALIARI, M., VIANELLO, M., DE MARCHI, S., AND MONTAGNA, R. 2006. HYPER2D: a numerical code for hyperinterpolation at Xu points on rectangles. *Appl. Math. Comp.* 183, 1138–1147.
- CARNICER, J. M., GASCA, M., AND SAUER, T. 2006. Interpolation lattices in several variables. *Numer. Math.* 102, 559–581.
- DE BOOR, C., DYN, N., AND RON, A. 2000. Polynomial interpolation to data on flats in R^d . *J. Approx. Theory* 105, 2, 313–343.
- DUNKL, C. F. AND XU, Y. 2001. *Orthogonal Polynomials of Several Variables*. Encyclopedia of Mathematics and its Applications, vol. 81. Cambridge University Press, Cambridge.
- FRANKE, R. 1982. Scattered data interpolation: tests of some methods. *Math. Comp.* 38, 181–200.
- GASCA, M. AND SAUER, T. 2000. Polynomial interpolation in several variables. *Adv. Comput. Math.* 12, 377–410.
- REIMER, M. 2003. *Multivariate Polynomial Approximation*. International Series of Numerical Mathematics, vol. 144. Birkhäuser, Basel.
- SAUER, T. 1995. Computational aspects of multivariate polynomial interpolation. *Adv. Comput. Math.* 3, 219–238.
- XU, Y. 1996. Lagrange interpolation on Chebyshev points of two variables. *J. Approx. Theory* 87, 220–238.