

APPROXIMATING OPTIMAL POINT CONFIGURATIONS FOR MULTIVARIATE POLYNOMIAL INTERPOLATION*

MARC VAN BAREL[†], MATTHIAS HUMET[†], AND LAURENT SORBER[†]

Dedicated to Lothar Reichel on the occasion of his 60th birthday

Abstract. Efficient and effective algorithms are designed to compute the coordinates of nearly optimal points for multivariate polynomial interpolation on a general geometry. “Nearly optimal” refers to the property that the set of points has a Lebesgue constant near to the minimal Lebesgue constant with respect to multivariate polynomial interpolation on a finite region. The proposed algorithms range from cheap ones that produce point configurations with a reasonably low Lebesgue constant, to more expensive ones that can find point configurations for several two-dimensional shapes which have the lowest Lebesgue constant in comparison to currently known results.

Key words. (nearly) optimal points, multivariate polynomial interpolation, Lebesgue constant, greedy add and update algorithms, weighted least squares, Vandermonde matrix, orthonormal basis

AMS subject classifications. 41A10, 65D05, 65D15, 65E05

1. Introduction. In several theoretical as well as computational mathematical problems, one wants to work with complicated multivariate functions. However, in a lot of cases performing operations with these original functions is cumbersome and requires an unacceptably high computational effort. A solution to this problem is to replace the original complicated function by a function that can be handled much more easily, e.g., polynomial functions. Within this space of simpler functions, we can look for the function that optimizes one of several possible criteria. One example is the minmax criterion, but the computational effort to find the function that minimizes the infinity norm error, is large. Instead an approximant can be found that is almost as good as the minmax approximant by interpolating the original function in certain well-chosen points. These points are chosen in an optimal or nearly optimal way with respect to minimizing the Lebesgue constant.

In this manuscript we develop several algorithms to compute point configurations for multivariate polynomial interpolation that have a low or even almost minimal Lebesgue constant for a given geometry. We will refer to them as “good” points and nearly optimal points, respectively. Interpolating in these points will yield good polynomial approximants for the geometry, compared to the minmax polynomial approximant.

For the problem of approximating univariate functions by polynomials in a typical compact set on the real line, i.e., an interval, both the theory and the corresponding software are well-developed. We refer to Chebfun, a MATLAB toolbox, whose theoretical foundation and several of its applications are described in the book by Trefethen [15]. If one transforms an arbitrary compact interval to the interval $[-1, 1]$, it turns out that different types of Chebyshev points not only form nearly optimal point configurations, but that the computation of the corresponding interpolant can be performed very efficiently (and accurately) by using the

*Received October 15, 2013. Accepted March 4, 2014. Published online on March 31, 2014. Recommended by Claude Brezinski. The research was partially supported by the Research Council KU Leuven, project OT/10/038 (Multi-parameter model order reduction and its applications), PF/10/002 Optimization in Engineering Centre (OPTEC), by the Fund for Scientific Research–Flanders (Belgium), G.0828.14N (Multivariate polynomial and rational interpolation and approximation), and by the Interuniversity Attraction Poles Programme, initiated by the Belgian State, Science Policy Office, Belgian Network DYSCO (Dynamical Systems, Control, and Optimization). Laurent Sorber is supported by a Flanders Institute of Science and Technology (IWT) doctoral scholarship.

[†]Numerical Approximation & Linear Algebra Group (NALAG) Department of Computer Science Katholieke Universiteit Leuven, Celestijnenlaan 200 A B-3001 Leuven, Belgium
(`{marc.vanbareel, mathias.humet, laurent.sorber}@cs.kuleuven.be`).

Fast Fourier Transform (FFT). The zero sets of other orthogonal polynomials, e.g., Legendre polynomials, have similar approximating properties but they can not be represented explicitly and the corresponding approximant cannot be computed equally efficient. For univariate rational interpolation, the so-called rational Chebyshev points are nearly optimal on the interval $[-1, 1]$ (see [17]).

The problem setting is more complicated in the multivariate case, because the geometry can take on more general forms (e.g., a polygon, a disk, . . .), in contrast to the univariate case where the typical geometry is the interval. Moreover the degree structure of the polynomial functions is more general. For a theoretical overview, we refer the interested reader to [1].

One of the criteria to determine the location of good points for polynomial approximation in a geometry, is minimizing the Lebesgue constant, which is the maximum of the Lebesgue function.¹ Points in some geometry are considered to be nearly optimal if the Lebesgue constant with respect to that geometry is small, and they are optimal if the Lebesgue constant is as small as possible. The Padua points seem to be the first known example of nearly optimal points for total degree polynomial interpolation in two variables, with a Lebesgue constant increasing like log square of the degree. The corresponding geometry is a square or a rectangle (or another derived form). These Padua points have been discovered and extensively studied by the Padova-Verona research group on “Constructive Approximation and Applications” (CAA-group) and their collaborators².

For other geometries there are no explicit representations known for (nearly) optimal points with respect to minimizing the Lebesgue constant. The CAA-group has developed MATLAB software to compute such nearly optimal points for several geometries, e.g., the disk and the simplex, not only for minimizing the Lebesgue constant but also for maximizing the corresponding Vandermonde determinant (Fekete-points) [7]. Initializing the software with reasonably nearly optimal points, it can also be used to derive point sets with a smaller Lebesgue constant than the initial set. A disadvantage of the software is that it is rather slow and therefore limited to a relatively small number of points.

In [3, 14] a faster, greedy algorithm is presented that uses built-in Matlab routines to compute QR or LU factorizations to compute approximate Fekete and Leja points. The underlying matrix is a Vandermonde matrix based on the total-degree product Chebyshev basis of the smallest rectangle containing the compact domain. The method works for “moderate” degrees.

On March 4, 2013, an extension of Chebfun was made available to work with functions in two variables defined on a rectangle³. The package provides very fast approximation algorithms by using FFT’s, but the geometry is limited to the rectangle.

In the recent paper [11] a method is developed to compute a “good” set of nodes for multivariate polynomial interpolation based on a greedy optimization algorithm. In each step of the greedy algorithm, a new node from a finite discretization of the domain Ω is added to the current set of “good” nodes. The properties of the method, i.e., the nodes are unstructured, the nodes are a sequence and are nested, and the geometries are arbitrary, are the same as for the greedy adding algorithm that is described in this paper. The resulting sets of interpolation nodes have “good” properties with respect to the value of the corresponding Lebesgue constant and Vandermonde determinant.

In this manuscript, we represent the polynomial functions using orthogonal bases with respect to a discrete inner product where the mass points are lying within the considered geometry. This leads to small condition numbers for the generalized Vandermonde matrices

¹The corresponding definitions are given in Section 2.

²<http://www.math.unipd.it/~marcov/CAA.html>

³<http://www2.maths.ox.ac.uk/chebfun/chebfun2/>

involved in the computations that allow us to find nearly optimal point configurations that are much larger compared to the point configurations obtained by currently known techniques.

Instead of solving the minmax problem (5.1), the algorithms of this manuscript tackle different, but related, optimization problems that approximately solve the same problem. Although the optima of these related problems do not coincide with the optima of the original minmax problem, they can be solved much more efficiently, making it possible to minimize the Lebesgue constant much more effectively. The first two algorithms of this manuscript use a greedy approach to find a set of “good” interpolation points for a general geometry. In contrast to existing methods, the greedy approach is not only used to generate a point set, adding points one by one, but also to update the resulting point set. The greedy add method is slower than the methods described in [3, 14], but this is due to the use of a more general basis than the monomial basis. The other algorithm described in this paper solves a non-linear weighted least squares optimization problem. By adapting the weights during several iterations, we obtain point configurations that are almost optimal.

The manuscript is divided into the following sections. In Section 2 the definition of the Lebesgue function and the Lebesgue constant is given. In Section 3, it is explained how a good approximation of the Lebesgue constant can be computed in an efficient way. Section 4 describes the representation that will be used for the multivariate polynomials given a certain geometry. Section 5 gives several algorithms to compute point configurations, ranging from cheap ones that produce non-optimal point configurations with a reasonably low Lebesgue constant, to more expensive ones that can find point configurations with an almost optimal Lebesgue constant. In Section 6 we show the results of applying these algorithms on several geometries for different degrees.

2. Lebesgue constant. Let Ω be a compact subset of \mathbb{R}^n . Consider the space \mathcal{P}_δ^n of polynomials in n variables having total degree $\leq \delta$.⁴ This space has dimension N with

$$(2.1) \quad N = \binom{\delta + n}{n}.$$

Consider a set $X = \{\mathbf{x}_k\}_1^N$ of N points in Ω and a basis $\{\phi_k\}_1^N$ for \mathcal{P}_δ^n . Let $V_X = [\phi_j(\mathbf{x}_i)]_{i,j}$ denote the generalized Vandermonde matrix for this basis in the points X . Given a function $f \in C(\Omega)$, we can approximate this function by computing the multivariate polynomial interpolant $p \in \mathcal{P}_\delta^n$ in the set of points X . Note that this interpolant is well defined and unique iff the generalized Vandermonde matrix V_X is nonsingular. If that is the case, the set of points X is called unisolvent for the space \mathcal{P}_δ^n .

DEFINITION 2.1 (Lebesgue function and Lebesgue constant). *Given a compact set $\Omega \subset \mathbb{R}^n$ and a set of points $X = \{\mathbf{x}_k\}_1^N \subset \Omega$ that is unisolvent for \mathcal{P}_δ^n . The Lebesgue function $\lambda_X(\mathbf{y})$ is defined as*

$$\lambda_X(\mathbf{y}) = \sum_{i=1}^N |l_i(\mathbf{y})|$$

with $l_i(\mathbf{y})$ the i th Lagrange polynomial, i.e.,

$$\begin{cases} l_i \in \mathcal{P}_\delta^n \\ l_i(\mathbf{x}_j) = \delta_{i,j}, & \text{for } i, j = 1, 2, \dots, N. \end{cases}$$

⁴ More general subsets of polynomials can be considered, i.e., having another degree structure in comparison to the total degree.

The Lebesgue constant Λ_X is defined as the maximum of the Lebesgue function $\lambda_X(\mathbf{y})$ for $\mathbf{y} \in \Omega$, i.e.,

$$\Lambda_X = \max_{\mathbf{y} \in \Omega} \lambda_X(\mathbf{y}).$$

The Lebesgue constant is a measure to compare the polynomial interpolant with the best polynomial approximant in the uniform norm. More precisely, for any function $f \in C(\Omega)$, let p denote the polynomial interpolant and p^* the best polynomial approximant in uniform norm, then

$$\|f - p\|_\infty \leq (1 + \Lambda_X) \|f - p^*\|_\infty.$$

Hence, when the Lebesgue constant Λ_X is small, we can find an approximation of a function f that is almost as good as the best polynomial approximation p^* , by just computing the polynomial interpolant p , which is generally much easier to compute than p^* .

The magnitude of the Lebesgue constant Λ_X depends heavily on the configuration of the points X in the compact subset Ω . Before we look for different algorithms to find point configurations with a low Lebesgue constant, the next section investigates how we can efficiently approximate the Lebesgue constant Λ_X .

3. Approximating the Lebesgue constant Λ_X . Computing the Lebesgue constant for a region $\Omega \subset \mathbb{R}^n$ is not an easy problem. Following the same approach as in [7], we approximate the Lebesgue constant by taking the maximum over a finite set $Y \subset \Omega$ of K well-chosen points

$$(3.1) \quad \Lambda_X \approx \max_{\mathbf{y} \in Y} \sum_{i=1}^N |l_i(\mathbf{y})|.$$

There are several possible candidates for the finite point set Y . We have chosen for point meshes generated by the package DistMesh [13], mainly because of its flexibility to create suitable meshes for many different geometries. As we explain in the following paragraphs, for many geometries like the square and the disk there are better meshes available, i.e., they give a better approximation of the maximum of a function with the same number of points. We emphasize that our algorithms can work with any choice of Y , and that our choice of DistMesh mostly provides a straightforward way to use fairly good meshes for any geometry.

In our context, a discretization Y of a domain Ω should have two important properties. The first is its quality of approximating the maximum of a function on the domain. The second is the fact that the mesh is used by the algorithms of Sections 5.1 and 5.2, where points of the output set X are extracted from the mesh. Since the output set should approximate an optimal point configuration (with minimal Lebesgue constant), and experimentally, optimal interpolation points are known to cluster near the boundary⁵, the mesh Y should be more dense near the boundary.

In what follows, first we briefly explain how DistMesh works. Then we discuss (Weakly) Admissible Meshes (WAM), why these meshes work well to approximate the maximum of a function and how DistMesh seems to be an AM as well. We also give some comments on the fact that the mesh is denser near the boundary. Finally, a numerical comparison of five different meshes is presented.

⁵We believe that this is true for convex geometries, but not for the “non-convex” part of a boundary, e.g., the non-convex part of the boundary of the L-shape. In Figure 6.2 we show a nearly optimal point configuration for the L-shape, that exhibits a low density of points near the non-convex part of the geometry.




FIG. 3.1. Example of a mesh generated by DistMesh for the L-shape consisting of 3475 points.

DistMesh [13] is a simple Matlab tool that generates unstructured triangular and tetrahedral meshes. The code is simple to use because the geometry is defined as a signed distance function, i.e., for each point this function returns the signed minimum distance between this point and the boundary of the geometry. The sign is negative inside the domain while it is positive outside the domain. The actual mesh generation uses the Delaunay triangulation routine in Matlab and tries to optimize the node locations by a force-based smoothing procedure. Using a weight function, the desired edge length distribution is specified by the user. When using DistMesh to generate a triangulation using a uniform weight function, it generates a triangular mesh where the lengths of all the edges are nearly equal as described in [13].

To give an idea of the meshes generated by DistMesh, Figure 3.1 shows a mesh for the L-shape consisting of 3475 points. In the examples of Section 6, we show the efficiency and effectiveness of using DistMesh to generate the set Y and give more details on the values of the parameters used in the numerical experiments.

Admissible meshes and weakly admissible meshes were introduced in [8] by Calvi and Levenberg as a tool to quantify the uniform approximation properties of discrete least squares polynomial approximation. Given a geometry Ω , an admissible mesh (AM) is a sequence of point sets $A(\delta)$ in function of the degree δ , that satisfies

$$(3.2) \quad \|p\|_{\Omega} \leq C(A(\delta), \Omega) \|p\|_{A(\delta)}, \quad p \in \mathcal{P}_{\delta}^n,$$

where for a set S , $\|p\|_S = \max_{x \in S} p(x)$ and where the constant $C(A(\delta), \Omega)$ is bounded above for all δ (see [8, (2.9)]). If the constant $C(A(\delta), \Omega)$ behaves like a polynomial in δ when $\delta \rightarrow \infty$, then the sequence of point sets $A(\delta)$ is called a weakly admissible mesh (WAM). Hence, if $C(A(\delta), \Omega)$ is small enough, (W)AMs are good discretizations of a geometry Ω to approximate the maximum of a polynomial of degree δ . In the numerical experiment described later, we indicate that point sets computed by DistMesh are WAMs.

The number of points K for a uniform AM behaves like $O(\delta^4) = O(N^2)$ when the degree δ goes to infinity. Since the number K of points increases very fast in function of the degree δ , for specific geometries AMs were constructed where K behaves as $O(N)$ [6]. See also [2, 4] on WAMs. These specific meshes have a higher density of points near the boundary.

Choosing point sets with more points in the neighborhood of the boundary is advantageous as can be seen as follows. When one has a nearly-optimal point set, e.g., on the unit square geometry, moving one of these points in the neighborhood of the boundary has a much larger influence on the Lebesgue function than moving a point in the center of the square. Hence, it seems better to increase the density of the points in the neighborhood of the boundary of the geometry. Taking the same number of points as for a uniform AM, this should not decrease the quality of the mesh, on the contrary.

In the following example, we compare five point sets on the unit square in \mathbb{R}^2 with respect to their quality as an AM. Three of the five point sets are generated by methods that can be used for general geometries: the uniform and non-uniform point sets generated by DistMesh, and a uniform covering of the unit square. The other two are specific AMs for the unit square: a non-uniform covering using Padua points and one using a tensor Chebyshev grid. To measure the quality of an AM ($A(\delta)$) the constant $C(A(\delta), \Omega)$ as defined in (3.2) can be estimated. The smaller this constant, the better. To compute a lower bound of $C(A(\delta), \Omega)$, we can rewrite (3.2) to get

$$C(A(\delta), \Omega) \geq \frac{\|p\|_{\Omega}}{\|p\|_{A(\delta)}}, \quad 0 \neq p \in \mathcal{P}_{\delta}^n.$$

We take 100 random polynomials p and use the maximum of all the fractions $\frac{\|p\|_{\Omega}}{\|p\|_{A(\delta)}}$ as a lower bound for $C(A(\delta), \Omega)$. The numerator is approximated by taking a finer discretization of Ω than $A(\delta)$. To compute an approximate upper bound, we use a similar method as described in [5]⁶. For given function values in each of the points of the point set $A(\delta)$, we consider the least squares approximating polynomial of degree δ . We approximate the maximum of the value in Ω of this approximating polynomial by taking the maximum value in a finer discretization of Ω than $A(\delta)$. The infinity norm of the operator going from the given function values to the function values on Ω gives an upper bound for $C(A(\delta), \Omega)$. By taking the finer discretization instead of Ω itself, an approximate upper bound is obtained.

We compare the values of the constant $C(A(\delta), \Omega)$ when Ω is the unit square in \mathbb{R}^2 . The lower and upper bounds for the constant $C(A(\delta), \Omega)$ are shown in function of δ in Figure 3.2. Each of the five point sets has approximately $8^4 = 4096$ points. The point set having a higher point density near the boundary generated by DistMesh works better than uniform point distributions having the same number of points. However, the specific admissible meshes (Padua points, Chebyshev tensor grid) developed for the geometry of the square are performing better to approximately maximize a given function. For a specific geometry having (weakly) admissible meshes (WAMs), it seems to be better to use such a WAM.

A detailed comparison between the different choices of the point sets and developing the corresponding theory is not within the scope of this paper.

⁶This method was suggested by one of the referees.

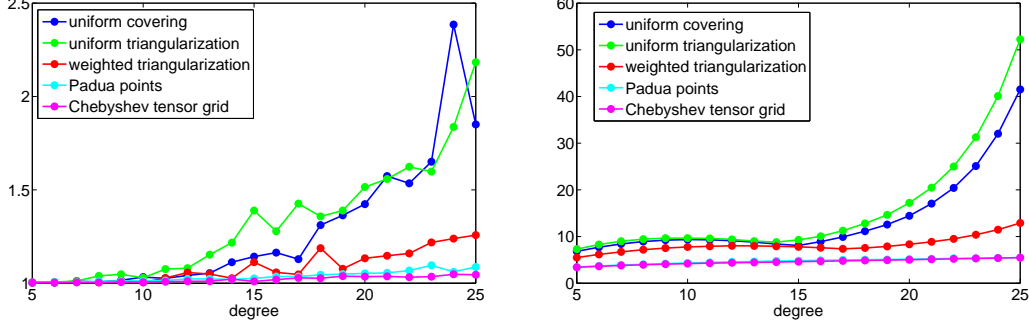


FIG. 3.2. Lower and upper bounds (left, respectively right figure) for $C(A(\delta), \Omega)$ in function of the degree δ

To compute the approximation (3.1), we choose a basis $\{\phi_k\}_1^N$ in \mathcal{P}_δ^n . More details on the choice of this basis will be given in Section 4. From the definition of Lagrange polynomials, we have the following expression for the basis polynomials:

$$[\phi_1(\mathbf{y}) \quad \cdots \quad \phi_N(\mathbf{y})] = [l_1(\mathbf{y}) \quad \cdots \quad l_N(\mathbf{y})] \begin{bmatrix} \phi_1(\mathbf{x}_1) & \cdots & \phi_N(\mathbf{x}_1) \\ \vdots & & \vdots \\ \phi_1(\mathbf{x}_N) & \cdots & \phi_N(\mathbf{x}_N) \end{bmatrix},$$

or evaluated in each of the K points $\mathbf{y}_j \in Y$:

$$\begin{bmatrix} \phi_1(\mathbf{y}_1) & \cdots & \phi_N(\mathbf{y}_1) \\ \vdots & & \vdots \\ \phi_1(\mathbf{y}_K) & \cdots & \phi_N(\mathbf{y}_K) \end{bmatrix} = \begin{bmatrix} l_1(\mathbf{y}_1) & \cdots & l_N(\mathbf{y}_1) \\ \vdots & & \vdots \\ l_1(\mathbf{y}_K) & \cdots & l_N(\mathbf{y}_K) \end{bmatrix} \begin{bmatrix} \phi_1(\mathbf{x}_1) & \cdots & \phi_N(\mathbf{x}_1) \\ \vdots & & \vdots \\ \phi_1(\mathbf{x}_N) & \cdots & \phi_N(\mathbf{x}_N) \end{bmatrix}.$$

We write this in a concise way as

$$(3.3) \quad V_Y = L V_X.$$

Note that K is chosen such that $K \gg N$.

The matrices V_X and V_Y are the basis polynomials evaluated in the points of the sets X and Y and V_X is the generalized Vandermonde matrix of the previous section. If the point set X is unisolvent, the matrix L of Lagrange polynomials can be computed by solving a system of linear equations with coefficient matrix V_X . Taking its matrix infinity norm results in approximation (3.1) of the Lebesgue constant, i.e.,

$$\Lambda_X \approx \|L\|_\infty = \|V_Y V_X^{-1}\|_\infty.$$

The accuracy of the computation of $\|L\|_\infty$ depends on the condition number of the generalized Vandermonde matrix V_X . For this number to be small, it is important to obtain a good basis $\{\phi_k\}_1^N$ for the geometry Ω considered, which we discuss in more detail in the next section.

4. Obtaining a good basis for a specific geometry. In this section we discuss some of the possible choices for the basis of \mathcal{P}_δ^n that are used to compute the Lebesgue constant Λ_X . First we mention the bases that have been used in [7] to obtain point configurations with a low Lebesgue constant for the square, the simplex and the disk. Then we discuss orthonormal

bases with respect to a discrete inner product, which can be computed by solving an inverse eigenvalue problem [16]. We briefly describe the problem setting and mention some of the approaches to solve the inverse eigenvalue problem. Finally we introduce a technique to extend a basis, which will be used in Section 5.1.

Since the choice of the basis determines the Vandermonde matrix V_X of the system (3.3), it has a large impact on the conditioning of the problem of computing Λ_X . The idea we pursue in this paper is to use a basis for which the condition number of V_X is small enough. The precise meaning of “small enough” depends on how accurate the computed value of Λ_X needs to be. For example, for the algorithms of Section 5, in practice it suffices to know only a couple of correct significant decimal digits of the matrix L in (3.3), so that $\text{cond}(V_X)$ may be as large as 10^{12} .

Briani et al. [7] use three different orthonormal bases for the respective geometries considered. Let $\Omega \in \mathbb{R}^n$ be a compact set, then we say that two polynomials $p, q \in \mathcal{P}_\delta^n$ are orthogonal with respect to Ω and the weight function $w(\mathbf{x})$ if

$$\langle p, q \rangle_\Omega := \int_\Omega p(\mathbf{x})q(\mathbf{x})w(\mathbf{x})d\mathbf{x} = 0.$$

The three bases consist of product Chebyshev polynomials for the square, Dubiner polynomials for the simplex and Koornwinder type II polynomials for the disk. These polynomials are orthonormal with respect to the respective geometries and the respective weight functions $w(\mathbf{x}) = \prod_{i=1}^n (1 - x_i)^{-\frac{1}{2}}$, $w(\mathbf{x}) = 1$ and $w(\mathbf{x}) = 1$.

Our approach is to consider a discrete inner product

$$(4.1) \quad \langle p, q \rangle_X = \sum_{i=1}^N w_i^2 p(\mathbf{x}_i)q(\mathbf{x}_i),$$

with points $X := \{\mathbf{x}_i\}_1^N \subset \mathbb{R}^n$ and weights $w_i \in \mathbb{R}^+$. An advantage of using an orthonormal basis $\{\phi_k\}_1^N$ with respect to this inner product is that, for $w_i = 1$, the matrix V_X is orthogonal. Hence, numerical difficulties to compute Λ_X for a set of points X can be avoided by taking an orthonormal basis with respect to (4.1) defined on the same point set X .

The problem of computing orthogonal multivariate polynomials with respect to (4.1) has been studied in [16]. In this work the orthogonal polynomials are represented by the recurrence coefficients $h_{i,j}^{(k)}$ of the recurrence relation

$$(4.2) \quad x_k \phi_j = \sum_{i=1}^{\pi_j^{(k)}} h_{i,j}^{(k)} \phi_i,$$

which gives an expression for $\phi_{\pi_j^{(k)}}$ if the previous polynomials $\phi_1, \dots, \phi_{\pi_j^{(k)}-1}$ are known.

The index $\pi_j^{(k)}$ depends on j and k and will be discussed later. The polynomials have to be ordered along a term order, meaning that $\phi_k(\mathbf{x}) = a_k \mathbf{x}^{\alpha_k} + \dots + a_1 \mathbf{x}^{\alpha_1}$ and the monomials $\mathbf{x}^{\alpha_k} := x_1^{\alpha_{k,1}} \dots x_n^{\alpha_{k,n}}$ satisfy a term order: $1 \prec \mathbf{x}^\beta$ for all $\beta \neq \mathbf{0}$ and if $\mathbf{x}^{\alpha_i} \prec \mathbf{x}^{\alpha_j}$, then $\mathbf{x}^\beta \mathbf{x}^{\alpha_i} \prec \mathbf{x}^\beta \mathbf{x}^{\alpha_j}$ for all $\beta \neq \mathbf{0}$. Here, we will restrict ourselves to graded term orders, imposing the additional condition that, if $\sum_k \alpha_{i,k} =: |\alpha_i| < |\alpha_j|$, then $\mathbf{x}^{\alpha_i} \prec \mathbf{x}^{\alpha_j}$. An example of a graded term order is the *graded lexicographical order*, which for $n = 3$ looks like

$$1 \prec z \prec y \prec x \prec z^2 \prec yz \prec y^2 \prec xz \prec xy \prec x^2 \prec \dots$$

A matrix expression for (4.2) is

$$x_k [\phi_1 \ \phi_2 \ \cdots \ \phi_{\hat{N}}] = [\phi_1 \ \phi_2 \ \cdots \ \phi_N] \hat{H}_k$$

with $\hat{H}_k(i, j) = h_{i,j}^{(k)}$ and $\hat{H}_k \in \mathbb{R}^{N \times \hat{N}}$. Here N and \hat{N} are the dimensions of the spaces \mathcal{P}_δ^n and $\mathcal{P}_{\delta-1}^n$, respectively. The element $h_{\pi_j^{(k)}, j}^{(k)}$ associated with the leading basis polynomial in (4.2) is called a pivot element of \hat{H}_k and it is the last nonzero element in the j -th column. The positions $(\pi_j^{(k)}, j)$ of the pivot elements follow from the monomial order and can be determined at a negligible cost. E.g., for the graded lexicographical ordering and $n = 3$, the matrix \hat{H}_x has pivots at positions

$$(4, 1), (8, 2), (9, 3), (10, 4), (15, 5), (16, 5), \dots$$

If $\mathbf{w} = [w_1 \ \dots \ w_N]^T$ is a vector with the weights and $X_k = \text{diag}(x_{1,k}, \dots, x_{N,k})$ is the diagonal matrix with the k -th coordinates of the points $\mathbf{x}_i \in X$, then the recurrence matrices \hat{H}_k can be found from the inverse eigenvalue problem

$$(4.3) \quad Q^T Q = I, \quad Q^T \mathbf{w} = \|\mathbf{w}\|_2 e_1, \quad \text{and} \quad H_k = Q^T X_k Q, \quad k = 1, \dots, n,$$

where the matrices \hat{H}_k are embedded in the $H_k \in \mathbb{R}^{N \times N}$ as follows

$$H_k = [\hat{H}_k \quad \times].$$

The basic idea is to apply orthogonal transformations to \mathbf{w} and X_k to make zeros in \mathbf{w} while at the same time assuring that the matrices H_k have the correct pivot element structure, which is determined by the monomial order. If the pivot elements in the matrices H_k are positive, then the process has a unique outcome.

We have implemented two methods to solve (4.3), where the user can supply any graded term order. The first method adds one points at a time. In each step, it uses Givens transformations to make one weight in \mathbf{w} zero and to bring the matrices H_k to the desired structure. The algorithm is explained in [16] for the bivariate case. The second method uses Householder transformations. A first Householder is applied to \mathbf{w} to make all the zeros at once. Subsequent Householders then bring H_k to the desired structure.

Although the method with Householder transformations has a higher flopcount than the method with Givens transformations, it becomes faster for large problems, because the operations are less granular. By using more matrix vector products instead of fine grain operations on vectors, most of the work is done using BLAS-2 routines (see [9, Chapter 1]). We will therefore prefer the second method for large problems, but the first method remains useful, because it allows to add points to an existing inner product.

As noted in [16], there is some freedom in the algorithms concerning which pivot is used to construct the Givens or Householder transformation. Several criteria to choose the pivot have been implemented, so the reader can experiment with them. We have adopted the approach to construct the orthogonal transformation from the vector with the highest 2-norm, since this seemed the most accurate in numerical tests. Numerical tests also pointed out to use a similar approach to evaluate the orthonormal polynomials using the recurrence relations (4.2): if $l = \pi_{j_1}^{(k_1)} = \dots = \pi_{j_m}^{(k_m)}$, so there are m pivot elements in the l -th row of respective matrices \hat{H}_{k_i} , then ϕ_l is computed from (4.2) for that k_i associated with the biggest pivot $h_{l,j_i}^{(k_i)}$.⁷

⁷Note that choosing the biggest pivot is similar to the optimal pivoting strategy for Gaussian elimination.

The last part of this section is devoted to explain a simple technique that extends a basis. In Section 5.1, we motivate this technique and give some numerical results that show its use. Suppose we have a basis $\{\phi_k\}_1^N$ for \mathcal{P}_δ^n associated with a graded term order, which is a good representation on a certain domain $\Omega \in \mathbb{R}^n$. We extend this basis with polynomials $\phi_{N+1}, \phi_{N+2}, \dots, \phi_{N+m}$ by taking products of the original basis

$$\phi_i = \phi_{k_i} \cdot \phi_{l_i}, \quad i = N + 1, \dots, N + m,$$

where the indices k_i and l_i satisfy

- (i) $\alpha_i = \alpha_{k_i} + \alpha_{l_i}$,
- (ii) $|\alpha_{l_i}| = |\alpha_{N+1}| - 1$,
- (iii) k_i is as low as possible.

Condition (i) follows directly from the definition of the monomial order and condition (ii) implies that we take the total degree of one of the factors to be one less than the total degree of the first polynomial that extends the basis. From (i) and (ii), the total degree of ϕ_{k_i} is fixed, and condition (iii) then determines the values of k_i and l_i .

Such an extension of a good basis on a domain will usually be less good than the original basis, and it is clear that it will deteriorate as m grows larger. However, the main advantage is that it can be evaluated very cheaply in points where the original basis has been evaluated. In Section 5.1 it is explained how this technique can be used to decrease computation time, while at the same time maintaining a high enough level of robustness.

5. Computing nearly optimal interpolation points. As explained in Section 2, we get a good polynomial approximation of the minmax polynomial approximant by interpolation in points X with a small Lebesgue constant Λ_X . To obtain such a set X , we want to solve the following minmax optimization problem

$$(5.1) \quad \min_{X \subset \Omega} \Lambda_X = \min_{X \subset \Omega} \max_{\mathbf{y} \in \Omega} \lambda_X(\mathbf{y}).$$

If we approximate the Lebesgue constant as in Section 3 by $\Lambda_X \approx \|L\|_\infty$, we get the optimization problem

$$(5.2) \quad \begin{aligned} & \min_{X \subset \Omega} \|L\|_\infty \\ & \text{subject to } V_Y = L V_X, \end{aligned}$$

where $X = \{\mathbf{x}_i\}_1^N$ and $Y = \{\mathbf{y}_i\}_1^K$.

This is a minmax optimization problem with constraints because the points \mathbf{x}_i have to lie in the region Ω . Minmax optimization problems are notoriously difficult to solve. In addition the objective function Λ_X is not everywhere differentiable, and the number of variables grows fast when increasing the degree δ and/or the number of dimensions n . E.g., for $n = 2$ and $\delta = 20$, the dimension N of the vector space \mathcal{P}_δ^n is 231. Hence, the number of real variables is the number of components of the N points \mathbf{x}_k , i.e., 462.

In [7], Briani et al. describe a collection of MATLAB scripts to solve the optimization problem (5.2) using the MATLAB Optimization Toolbox. They consider $n = 2$ and Ω equal to the square, the disk and the simplex, and their results include nearly optimal point configurations for these geometries up to a total degree of $\delta = 20$. There is no certainty that the real optimum is reached, but the Lebesgue constants found are the smallest at the point of their writing.

In the next subsections, we present alternative methods to find a point set X with a low Lebesgue constant. These methods work for very general geometries, can be used for

larger point sets and are faster compared to current techniques. The first algorithm uses a relaxed optimization criterion and creates a point configuration with a relatively low Lebesgue constant in an efficient, non-iterative way. The second algorithm iterates over the point set one point at a time, using the same criterion. The third and fourth algorithm are more advanced optimization algorithms that solve a similar but easier problem than (5.2) leading to point configurations with a nearly optimal Lebesgue constant.

5.1. Greedy algorithm by adding points. Evaluating the objective function $\|L\|_\infty$ of the optimization algorithm (5.2) requires the evaluation of the basis in the points X and the solution of a system of linear equations. Since the objective function is not differentiable on Ω and the number of variables can become very high, the convergence to a local minimum using standard MATLAB Optimization tools can take a lot of iterations, and consequently a lot of objective function evaluations.

In this section, we develop a “greedy” algorithm to generate a point configuration for any geometry Ω with a reasonably low Lebesgue constant Λ_X , with only a small computational effort. The algorithm is based on two ideas:

1. In each step, one point of the region Ω is added, while the other points remain where they are.
2. This point is added there where the Lebesgue function reaches its maximum.

We will refer to the algorithm as the Greedy Add algorithm.

Criterion 2 is reasonable in the sense that it guarantees that the updated Lebesgue constant has the value 1 in the new point. This point can be approximated by taking it from the set $Y \subset \Omega$, where the Lebesgue function reaches a maximum. Note that the new point could also be chosen to minimize the Lebesgue constant as a function of only one point, but this would be much more costly. Instead, we use a greedy approach where the next point is picked based on the mentioned relaxed criterion. Numerical experiments will show that, although the point configurations obtained are clearly not optimal, they exhibit a structure in the domain Ω similar to (nearly) optimal configurations, and their Lebesgue constant is reasonably low.

Obviously, the first point can be chosen freely. Since the Lebesgue function for one point is a constant, the same holds for the second point. However, one must be careful to keep the set of two points unisolvent. E.g., consider the term order $1 \prec x \prec y \prec x^2 \prec \dots$. The Vandermonde matrix for two points (x_1, y_1) and (x_2, y_2) is

$$\begin{bmatrix} 1 & x_1 \\ 1 & x_2 \end{bmatrix}.$$

Hence the first two points can be chosen freely, but they must have different first coordinates.

Theoretically, it is possible that at some step, after adding the next point, the configuration is not unisolvent anymore. As a result, the Lebesgue constant reaches infinity, leaving the next point undefined. We give an example where Ω is the unit disk, with the same term order as just described. Suppose that the first two points are $(-1, 0)$ and $(0, 0)$. The Lebesgue function depends on x only and a small calculation shows that it reaches its maximum in the disk at $(1, 0)$. If this point is included in the point set Y , then the resulting point configuration of 3 points will not be unisolvent causing the method to fail. Since all points are collinear, the Vandermonde matrix is singular. If both points are chosen randomly, we believe that the probability for such an event to occur is zero.

Suppose that we want to generate a configuration of N points, where N is the dimension (2.1) of the space \mathcal{P}_δ^n .⁸ For now assume that there is a suitable basis for \mathcal{P}_δ^n on the geometry Ω , e.g., product Chebyshev polynomials on the square $[-1, 1]^2$. If that is the case, the Greedy Add algorithm can be formulated as Algorithm 1. Since the grid Y consist of K points, the matrix V_Y is of dimension $K \times N$. Furthermore, $V_Y^{(k)}$ is the $K \times k$ matrix with the first k columns of V_Y and $V_X^{(k)}$ is the $k \times k$ (generalized) Vandermonde matrix for the first k basis polynomials and the points in X . In step k , the matrix L is $K \times (k - 1)$ and each columns contains one of the Lagrange polynomials for the points in X evaluated in Y . The index i selects the point in Y where the Lebesgue function is maximal.

Algorithm 1 Greedy Add algorithm

Input: N, Y , basis

Output: X

```

 $X \leftarrow \{2 \text{ random points } \mathbf{x}_1 \text{ and } \mathbf{x}_2\}$ 
 $V_Y \leftarrow \text{evaluate basis functions in grid } Y \in \Omega$ 
for  $k = 3, \dots, N$  do
   $V_X^{(k-1)} \leftarrow \text{evaluate basis functions in } X$ 
   $L \leftarrow V_Y^{(k-1)} = LV_X^{(k-1)}$ 
   $i \leftarrow \text{index of row of } L \text{ with largest one norm}$ 
   $\mathbf{x}_k \leftarrow Y(i)$ 
   $X \leftarrow X \cup \{\mathbf{x}_k\}$ 
end for

```

Two remarks have to be made. First, the computation of L can be accelerated using the Sherman-Morrison-Woodbury formula ([9, p. 50]). Indeed, in step k the matrices $V_X^{(k-1)}$ and $V_Y^{(k-1)}$ are the same as in the previous step, except for the their last columns and the last row of $V_X^{(k-1)}$. The matrix of the system is therefore a rank-2 update of the system in the previous step. Making use of this fact improves the efficiency of one step from $O(Kk^2)$ flop to $O(Kk)$. There should be a $O(k^3)$ term as well, but we get rid of it by updating the QR factorization of $V_X^{(k-1)}$ ([9, Section 12.5]).

Second, given a geometry Ω , it is not always apparent which basis to use, if the Vandermonde matrices in the algorithm have to remain well conditioned. As an example, we carry out Algorithm 1 on the L-shape using product Chebyshev polynomials for a degree $\delta = 30$ or $N = 496$, and we plot the condition number of $V_X^{(k)}$ in Figure 5.1. The condition number keeps growing steadily until at some point it becomes so large that the Lebesgue function evaluations possibly have no correct significant digits left.

A solution to this problem is using polynomials orthogonal with respect to a discrete inner product (4.1) with the current points in step k . In this way, the Vandermonde matrix is always perfectly conditioned. This solution involves solving the inverse eigenvalue problem (4.3) of size k in every step, after finding the next point, and evaluating the new set of orthogonal polynomials in the points Y . The inverse eigenvalue problem can be updated one point at the time using the Givens implementation (see Section 4), at a cost of $O(k^2)$ flops per step. Hence, the expensive part of the process is evaluating the new basis functions in the points Y at a cost of $O(Kk^2)$ flops per step.

To avoid the costly procedure of updating the basis in each step, we try to extend the

⁸Note that all the algorithms work for any value of N , but for notational convenience we work with spaces of total degree.




FIG. 5.1. The condition number of the Vandermonde matrix $V_X^{(k)}$ using Algorithm 1 for the L-shape $\Omega = [-1, 1] \times [-1, 0] \cap [-1, 0] \times [0, 1]$, with product Chebyshev polynomials as basis.

current basis with products of the original basis functions, as explained in Section 4. We keep track of the reciprocal condition number of $V_X^{(k)}$, which is cheap to compute⁹, and only if $V_X^{(k)}$ becomes too badly conditioned we compute a new orthogonal basis. In Figure 5.2 we plot again the condition number of $V_X^{(k)}$ for the L-shape, now using the adaptations just described. The condition number grows steadily, but once it becomes too large, the basis is updated. For $N = 496$, only 2 costly basis updates have been carried out, which is a significant improvement.

Each time the basis is updated, we recompute the matrix L by solving a regular linear system. Note that this is not strictly necessary, since the Lagrange polynomials are independent of the basis that is used, so it is possible to continue updating L via low rank updates. However, it might be useful to avoid inaccuracies in the matrix L obtained by the subsequent low rank updates. A stability analysis of these updates is not covered in this paper.

Since the implementation of the adapted Greedy Add Algorithm is a bit too technical to be included in this paper, we refer to the documentation in the code. In Figure 5.3 the value of the Lebesgue constant is plotted for each iteration of the adapted algorithm, for several pairs of random starting points and for several sizes of the grid Y . Observe that the Lebesgue constant fluctuates a lot, and that the final value Λ_N can be a lot larger than the previous value. This shows that the obtained point configurations are by no means optimal, but they can serve as a starting point for the algorithms in the following sections. In addition, observe that the choice of the starting points influences the obtained Lebesgue constants, as does the size of the grid Y .

The resulting point configuration is shown in Figure 5.4 for one particular choice of the starting points and the size of the grid, for both the square and the L-shape. In Section 6 we obtain point configurations with nearly optimal Lebesgue constants, which are shown in Figure 6.2. We observe that the structure in these optimal point configurations is already present in the point configurations obtained by the Greedy Add Algorithm.

5.2. Greedy algorithm by updating points. In this section we develop the Greedy Update Algorithm, implementing a straightforward approach to improve the point configuration $X = \{x_k\}_1^N$ obtained by the Greedy Add Algorithm of the previous section. The idea is iter-

⁹MATLAB's RCOND gives an approximation of the reciprocal condition number $\text{cond}(V_X^{(k)})^{-1}$.




FIG. 5.2. The condition number of the Vandermonde matrix $V_X^{(k)}$ using the adapted version of Algorithm 1 for the L-shape, with orthogonal polynomial updates and basis extension.




FIG. 5.3. The Lebesgue constant Λ_k after adding the first k points with the adapted Greedy Add Algorithm as a function of k , for several random choices of the first two points (left) and for several sizes of the grid Y (right). The geometry Ω is the square and $\delta = 20$, so $N = 231$. The gridsize for the left plot is 21322.

ate over all the points, remove each point from X and immediately add a new point according to the same greedy criterion. By iterating several times over all the points, the Lebesgue constant typically stabilizes at a reasonably low value.

The algorithm is described schematically in Algorithm 2. The input variables are a point configuration X , e.g., obtained by the Greedy Add Algorithm, a grid $Y \in \Omega$ and the variables needed to evaluate the basis that is used. One possibility is an basis orthogonal with respect to X . We have observed that if the input point configuration X has a low enough Lebesgue constant, then this basis will remain good enough for all the iterations. We have added the functionality that the basis is updated if the Vandermonde matrix $V_X^{(N-1)}$ becomes too badly conditioned.

Similar to the Greedy Add Algorithm, the computation of L in each step can be accelerated by using low rank updates. Indeed, the matrix $V_X^{(N-1)}$ in step $k + 1$ is identical to $V_X^{(N-1)}$ in step k , except for its k -th row. They are the same basis polynomials (columns) evaluated in the same points (rows) except for one. Hence, the matrix of the system is a rank-1 update of the system in the previous step and we can again reduce the amount of work in one step from $O(KN^2)$ to $O(KN)$ flop. The QR factorization of V_X^{N-1} is updated as well.




FIG. 5.4. Point configurations X of $N = 231$ points ($\delta = 20$) obtained by the Greedy Add Algorithm for the square on the left, and the L-shape on the right.

Algorithm 2 Greedy Update Algorithm

Input: X, Y , basis

Output: X

```

 $V_Y \leftarrow$  evaluate basis functions in grid  $Y \in \Omega$ 
while stopping criterion is not satisfied do
  for  $k = 1, 2, \dots, N$  do
     $X \leftarrow X \setminus \{\mathbf{x}_k\}$ 
     $V_X^{(N-1)} \leftarrow$  evaluate basis functions in  $X$ 
     $L \leftarrow V_Y^{(N-1)} = LV_X^{(N-1)}$ 
     $i \leftarrow$  index of row of  $L$  with largest one norm
     $\mathbf{x}_k \leftarrow Y(i)$ 
     $X \leftarrow X \cup \{\mathbf{x}_k\}$ 
  end for
end while

```

Figure 5.5 is an extension of Figure 5.3, where the value of the Lebesgue constant is plotted for each iteration of the adapted Greedy Add Algorithm and the Greedy Update Algorithm, for several pairs of random starting points and for several sizes of the grid Y . The Greedy Update Algorithm runs for 10 iterations over all the points. We observe that usually the Lebesgue constant stabilizes after a couple of runs and that the value of the final Lebesgue constant depends on the particular choice of the starting points and on the size of the grid.

5.3. Algorithm based on approximating the infinity norm. The infinity norm in (5.2) is notoriously difficult to optimize using numerical optimization techniques because it combines two of the most exacting objective function properties: taking the maximum over a set and summing (nonsmooth) absolute values. For many initial point sets X , the Lebesgue constant will be quite large and it may suffice to solve a neighbouring problem approximating (5.2) in order to obtain a substantial reduction of the Lebesgue constant.

5.3.1. Unweighted least squares problem. One approach could be to replace the infinity norm by the (squared) Frobenius norm since

$$(5.3) \quad \frac{1}{\sqrt{KN}} \|L\|_F \leq \|L\|_\infty \leq \sqrt{N} \|L\|_F.$$




FIG. 5.5. The Lebesgue constant Λ_k at each iteration of the adapted Greedy Add Algorithm and the Greedy Update Algorithm, for several random choices of the first two points (left) and for several sizes of the grid Y (right). The geometry Ω is the square and $\delta = 20$, so $N = 231$. The gridsize for the left plot is 21322.

For example, for $n = 2$ and $\delta = 20$ we have $N = 231$ and hence $\|L\|_F$ bounds $\|L\|_\infty$ from above by about a factor of 15. In practice, the two norms are often even closer than the bound (5.3) suggests. The objective is now to solve the optimization problem

$$(5.4) \quad \min_{X \subset \Omega} \frac{1}{2} \|L\|_F^2$$

subject to $V_Y = L V_X$.

By eliminating the (linear) constraint in (5.4), we obtain a nonlinear least squares (NLS) problem in $X \subset \Omega$. There are several algorithms for solving NLS problems, many of which can be adapted for solutions restricted to a domain Ω . In our experiments, we use a projected Gauss-Newton dogleg trust-region method, which is a straightforward generalization of the bound-constrained projected Newton algorithm of [10] to a larger class of geometries. To define a geometry Ω , the user is asked to implement a function which projects points outside of the geometry onto its boundary.

Given a current iterate, the Gauss-Newton dogleg trust-region method computes two additive steps. The first is the Cauchy step p_{CP} , which is approximated as a scaled steepest descent direction $-g := \frac{df(z)}{dz}$, where $z := \text{vec}(X)$ ¹⁰ and the objective function $f(z)$ is defined as $\frac{1}{2} \|L\|_F^2$. Here, the points X are stored as $[x_i^{(j)}]_{i,j}$, where $x_i^{(j)}$ is the j th component of the i th point. The second is the Gauss-Newton step

$$(5.5) \quad p_{GN} := -\text{red}(J^T J)^\dagger g,$$

where the Jacobian J is defined as $\frac{d\text{vec}(L)}{dz^T}$, and $\text{red}(\cdot)$ “reduces” the Hessian approximation $J^T J$ by setting those rows and columns corresponding to the active set equal to those of the identity matrix of the same size as $J^T J$. The active set is defined as the set of indices i for which the variables z_i are on the boundary of the geometry. For more details on the reduction of the Hessian; see [10]. Since J is tall and skinny, its Gramian $J^T J$ is a relatively small square matrix of order Nn . Furthermore, it is a positive (semi-)definite approximation of the objective function’s Hessian and hence may be expected to deliver a high-quality descent direction p_{GN} for a relatively low computational cost. Importantly, we will see that computing the two descent directions can be done with an amount of computational effort that is independent of the number of mesh points K .

¹⁰If X is stored in MATLAB as a $N \times n$ matrix, then $\text{vec}(X) := X(:)$ is the $Nn \times 1$ vectorization of X .

To compute the aforementioned descent directions, let

$$W^{(i)} := \left[\frac{\partial(V_X)_1^T}{\partial x_1^{(i)}} \quad \dots \quad \frac{\partial(V_X)_N^T}{\partial x_N^{(i)}} \right]^T$$

be a compact representation of the derivative of V_X with respect to the i th component of the points X . Furthermore, let

$$W := [W^{(1)T} \quad \dots \quad W^{(N)T}]^T,$$

then after some straightforward computation we find that

$$-\mathbf{g} = -J^T \text{vec}(L) = [(1_{n \times 1} \otimes (V_X^{-T} (V_Y^T V_Y) V_X^{-1})) * W] 1_{N \times 1}$$

and

$$J^T J = (1_{n \times n} \otimes (V_X^{-T} (V_Y^T V_Y) V_X^{-1})) * (W V_X^{-1} V_X^{-T} W^T),$$

where $1_{m \times n}$ is an $m \times n$ matrix of ones, \otimes and $*$ are the Kronecker and Hadamard (or elementwise) product, respectively. Notice that the only computation involving vectors of length K is the term $V_Y^T V_Y$, which need only be computed once and can be done on beforehand. Consequently, the cost per Gauss–Newton iteration is dominated by the cost of solving (5.5), which requires $O(N^3 n^3)$ flop.

Once the Cauchy and the Gauss–Newton steps are computed, the projected Gauss–Newton dogleg trust-region algorithm proceeds to project them in such a way that the sum of the current iterate \mathbf{z}_k and these steps does not exceed the boundary of the geometry. In other words, using the user-defined projection function $\text{proj}(\cdot)$, the steps are corrected as

$$\mathbf{p} \leftarrow \text{proj}(\mathbf{z}_k + \mathbf{p}) - \mathbf{z}_k.$$

The dogleg trust-region algorithm then searches for a step which improves the objective function in (a subspace of) the plane spanned by the projected Cauchy and Gauss–Newton steps. For more details on dogleg trust-region; see, e.g., [12].

5.3.2. Weighted least squares problem. Because the Frobenius norm is only a crude approximation for the infinity norm, we introduce a diagonal weighting matrix $D_w = \text{diag}(d_w(i))$ in the least squares optimization problem (5.4):

$$(5.6) \quad \min_{X \in \Omega} \frac{1}{2} \|D_w L\|_F^2$$

subject to $V_Y = L V_X$.

This problem is solved in an approximate way by performing a small number of Gauss–Newton dogleg trust-region iteration steps¹¹. Based on this new approximate solution, the weights $d_w(i)$ are adapted. More weight is put on the points $\mathbf{y}_i \in Y \subset \Omega$ where the Lebesgue function is large. Solving the least squares problem with the adapted weights (5.6), generically pushes the Lebesgue function down in those subregions where more weight was placed.

To obtain an efficient and effective algorithm, it is crucial to design a good heuristic for this adaptation of the weights. By trial and error, the following heuristic came out as a good choice and was implemented. The number of points \mathbf{y}_i of the set Y is chosen approximately

¹¹In our implementation, the number of iterations is taken equal to two.

equal to one hundred times the number of points x_i of X . In total there are one hundred outer iterations each with another adapted weight matrix D_w . Initially the weights are all equal to one. After each outer iteration k the Lebesgue function is computed in all points y_i and the first $n_y(k)$ largest values are considered. The weight of each of the corresponding points is increased by a fixed amount δ_w , taken equal to 0.4 in our implementation. Note that the number $n_y(k)$ of points y_i whose weight is increased, depends on the index of the outer iteration. The formula for this number is

$$n_y(k) = \max\{10, N - \lfloor \frac{N}{60}k \rfloor\}$$

with $\lfloor r \rfloor$ the largest integer number less than or equal to the real number r . Hence, in each subsequent iteration, less points are receiving a higher weight until this number is equal to 10 after which it remains constant.

6. Numerical experiments. The algorithms were implemented in MATLAB R2012a and can be obtained from the corresponding author. The experiments were executed on a Linux machine with 2 Intel Xeon Processors E5645 and 48 GByte of RAM.

6.1. Experiment 1: nearly-optimal point configurations for the square, simplex, disk and L-shape. For each of the geometries, the square, simplex, disk and L-shape, a nearly optimal point configuration X is computed for each of the total degrees $\delta = 3, 4, \dots, 30$. To derive these points, the different optimization algorithms of Section 5 are used subsequently.

First, the Greedy Add Algorithm of Section 5.1 is used to obtain an initial configuration X_1 with a reasonably small Lebesgue constant. The point set Y_1 from which these initial points are taken, is generated by DistMesh with the parameter determined such that approximately $100N$ points are contained in set the Y_1 where N is the number of points of X_1 . This initial configuration is then improved by performing 2 iterations of the Greedy Update Algorithm of Section 5.2, using the same points Y_1 as in the first phase. This improved point configuration X_2 is the initialization of the final phase where the weighted least squares optimization algorithm from Section 5.3.2 is used. For the disk, the same point set Y_1 is used in this final phase. For the polygon-geometries, we generate a triangular mesh based on the points of X_2 together with the edge points of the polygon (square, simplex, L-shape). Each triangle is then divided in a number of subtriangles such that the side lengths are 10 times smaller. This results in a point set Y_2 that contains approximately $100N$ points. Performing 100 outer iterations of the weighted least squares algorithm results in the nearly optimal point configuration $X = X_3$.

In Figure 6.1 the estimated Lebesgue constant of the resulting set $X = X_3$ is shown for the square, simplex, L-shape and disk, respectively. The estimation of the Lebesgue constant is done by sampling the Lebesgue function on a point set generated as Y_2 based on the point set X_3 for degree 30 and with a multiplication factor 1000 instead of 100. In the subfigures also the results obtained by the CAA-group [7] are given when available.

The optimization problem that is solved, is not an easy one. For some degrees the computed minimal Lebesgue constant appears closer to the global optimum than for others. Like for many optimization problems, it is often difficult to reach global optimality. The plots show that it seems more difficult to compute nearly-optimal point sets for the simplex and L-shape than for the square and the disk.

These values of the minimal Lebesgue constant are the best that have been computed so far. For the square a non smooth behavior of the results of the CAA-group is clear from the

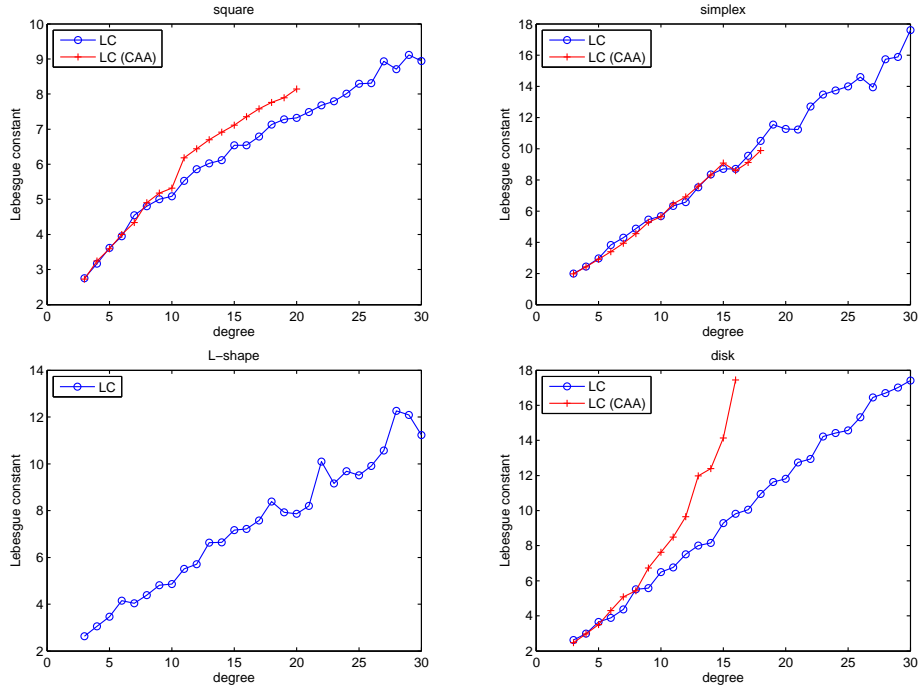


FIG. 6.1. Lebesgue constant in function of the degree for the square, simplex, L-shape and disk




FIG. 6.2. Nearly optimal point configurations of degree 30 for the square, simplex, L-shape and disk

graph¹². The graph showing our results for the square has a smoother behavior and for the disk it is clear that the growth as indicated by the results of the CAA group is a side-effect of their optimization procedure and not the real growth of the (nearly-)optimal Lebesgue constant.

In Figure 6.2 the corresponding nearly optimal point configurations for degree 30 are given. We have no formal way to compare the computed points with optimal ones because these optimal sets are not known. However, the shape of the Lebesgue function gives a good indication of the quality of the “good” points. When solving the minmax optimization problem, the points are subsequently moved in order to lower the highest local maximum of the Lebesgue function, while at the same time trying to keep the Lebesgue function low enough everywhere else. Intuitively, we expect that for the optimal point set, all the local maxima of the Lebesgue function will have the same function value. Conversely, when all the local maxima of the Lebesgue function have (almost) the same function value, the corresponding point set could be considered as a (nearly-)optimal one.

In Figure 6.3 the Lebesgue function for the Padua points (Lebesgue constant ≈ 9.2) and the Lebesgue function for “our” nearly-optimal point configuration (Lebesgue constant ≈ 7.3) on the square for degree 20 is shown indicating the difference in behavior of the local maxima. Note that the figure for the Padua points shows that the density of these points is not high enough in the neighborhood of the boundary. In Figure 6.4 the time for each of the three phases of the algorithm is presented. The lower curve is the time (in seconds) in function of the degree for the Greedy Add Algorithm. The middle curve shows the time for the Greedy Update Algorithm. The upper curve presents the time for the weighted least squares phase.

Compared to the algorithms of [7], to obtain a comparable Lebesgue constant the combined algorithm of this paper needs less computing time.

6.2. Experiment 2: nearly optimal point set for degree 60 on the square. This experiment shows that much larger nearly optimal point sets can be generated compared to existing techniques. For degree $\delta = 60$, the number of points is $N = 1891$ which is more than 8 times larger than for degree $\delta = 20$. For this experiment, we run only the two first phases of our combined optimization scheme, i.e., greedy adding and greedy updating, with 10 instead of 2 iterations for the greedy update step. The greedy add step takes 2.16 hours, while the greedy update step takes 23.38 hours. In Figure 6.5, the estimated Lebesgue constant is shown for each of the 10 iterations of the greedy update step as well as the resulting nearly optimal point configuration having an estimated Lebesgue constant of 75 which was reached in iteration 5.

7. Conclusion. In this paper several optimization algorithms were designed to compute nearly optimal point configurations for different geometries. These algorithms can be combined to derive an efficient and effective algorithm where one algorithm uses the output of the previous one as an initialization. By choosing a representation of the multivariate polynomials in terms of an orthogonal basis with respect to a discrete inner product for a geometry, numerical problems are avoided for larger point sets. Also the efficiency is at least one order of magnitude better compared to existing techniques.

In future research several topics can be studied:

- The different algorithms of Section 5 can be combined in many ways with different heuristics for the number of iterations in the greedy algorithm for updating and the inner and outer iteration of the weighted least squares algorithm. Also different point sets Y can be used in each of the algorithms.

¹²We have asked the authors of the corresponding paper if there had occurred a typo in presenting their results, but this didn't seem to be the case.




FIG. 6.3. Comparison of the Lebesgue function for the square and degree 20 for the Padua points (above) and the nearly-optimal point set that we found (below)

- It is not clear to us if the weighted least squares algorithm that has been developed to approximately solve the minmax optimization problem is known in the literature. At this point it uses a crude heuristic and more investigation is necessary to make this approach more robust. The generalization of this approach to other minmax optimization problems can be studied.




FIG. 6.4. Time (in seconds) for the three phases (greedy add, greedy update, weighted nonlinear least squares) of the algorithm for the square




FIG. 6.5. left: Lebesgue constant after each of the 10 iterations of the greedy update step for degree 60 on the square; right: resulting nearly optimal point configuration

REFERENCES

- [1] T. BLOOM, L. BOS, J.-P. CALVI, AND N. LEVENBERG, *Polynomial interpolation and approximation in C^d* , Ann. Polon. Math., 106 (2012), pp. 53–81.
- [2] L. BOS, J.-P. CALVI, N. LEVENBERG, A. SOMMARIVA, AND M. VIANELLO, *Geometric weakly admissible meshes, discrete least squares approximations and approximate Fekete points*, Math. Comp., 80 (2011), pp. 1623–1638.
- [3] L. BOS, S. DE MARCHI, A. SOMMARIVA, AND M. VIANELLO, *Computing multivariate Fekete and Leja points by numerical linear algebra*, SIAM J. Numer. Anal., 48 (2010), pp. 1984–1999.
- [4] ———, *Weakly admissible meshes and discrete extremal sets*, Numer. Math. Theory Methods Appl., 4 (2011), pp. 1–12.

- [5] L. BOS, A. SOMMARIVA, AND M. VIANELLO, *Least-squares polynomial approximation on weakly admissible meshes: disk and triangle*, J. Comput. Appl. Math., 235 (2010), pp. 660–668.
- [6] L. BOS AND M. VIANELLO, *Low cardinality admissible meshes on quadrangles, triangles and disks*, Math. Inequal. Appl., 15 (2012), pp. 229–235.
- [7] M. BRIANI, A. SOMMARIVA, AND M. VIANELLO, *Computing Fekete and Lebesgue points: Simplex, square, disk*, J. Comput. Appl. Math., 236 (2012), pp. 2477–2486.
- [8] J.-P. CALVI AND N. LEVENBERG, *Uniform approximation by discrete least squares polynomials*, J. Approx. Theory, 152 (2008), pp. 82–100.
- [9] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, 3rd ed., Johns Hopkins University Press, Baltimore, 1996.
- [10] C. T. KELLEY, *Iterative Methods for Optimization*, SIAM, Philadelphia, 1999.
- [11] A. NARAYAN AND D. XIU, *Constructing nested nodal sets for multivariate polynomial interpolation*, SIAM J. Sci. Comput., 35 (2013), pp. 2293–2315.
- [12] J. NOCEDAL AND S. J. WRIGHT, *Numerical Optimization*, 2nd ed., Springer, New York, 2006.
- [13] P.-O. PERSSON AND G. STRANG, *A simple mesh generator in MATLAB*, SIAM Rev., 46 (2004), pp. 329–345.
- [14] A. SOMMARIVA AND M. VIANELLO, *Computing approximate Fekete points by QR factorizations of Vandermonde matrices*, Comput. Math. Appl., 57 (2009), pp. 1324–1336.
- [15] L. N. TREFETHEN, *Approximation Theory and Approximation Practice*, SIAM, Philadelphia, 2012.
- [16] M. VAN BAREL AND A. A. CHESNOKOV, *A method to compute recurrence relation coefficients for bivariate orthogonal polynomials by unitary matrix transformations*, Numer. Algorithms, 55 (2010), pp. 383–402.
- [17] J. VAN DEUN, K. DECKERS, A. BULTHEEL, AND J. A. C. WEIDEMAN, *Algorithm 882: Near-best fixed pole rational interpolation with applications in spectral methods*, ACM Trans. Math. Software, 35 (2008), Art. 14, (21 pages).