Ergod. Th. & Dynam. Sys. (2005), **25**, 321–352
doi:10.1017/S0143385704000537

Convex dynamics and applications

R. L. ADLER†, B. KITCHENS‡, M. MARTENS§, C. PUGH¶, M. SHUB†‖
and C. TRESSER†

† *IBM, TJ Watson Research Center, Yorktown Heights, NY 10598-0218, USA*
(e-mail: rla@us.ibm.com, shub@math.toronto.edu, tresser@us.ibm.com)
‡ *Mathematical Sciences Department, IUPUI-LD270, Indianapolis, IN 46202, USA*
(e-mail: bkitchens@math.iupui.edu)
§ *University of Groningen, Department of Mathematics, PO Box 800,*
9700 AV Groningen, The Netherlands
(e-mail: marco@math.rug.nl)
¶ *Mathematics Department, University of California, Berkeley, CA 94720, USA*
(e-mail: pugh@math.berkeley.edu)

(*Received* 4 *March* 2004 *and accepted in revised form* 10 *June* 2004)

Abstract. This paper proves a theorem about bounding orbits of a time dependent dynamical system. The maps that are involved are examples in convex dynamics, by which we mean the dynamics of piecewise isometries where the pieces are convex. The theorem came to the attention of the authors in connection with the problem of digital halftoning. *Digital halftoning* is a family of printing technologies for getting full-color images from only a few different colors deposited at dots all of the same size. The simplest version consists in obtaining gray-scale images from only black and white dots. A corollary of the theorem is that for *error diffusion*, one of the methods of digital halftoning, averages of colors of the printed dots converge to averages of the colors taken from the same dots of the actual images. Digital printing is a special case of a much wider class of scheduling problems to which the theorem applies. Convex dynamics has roots in classical areas of mathematics such as symbolic dynamics, Diophantine approximation, and the theory of uniform distributions.

1. *Introduction*

THEOREM 1.1. *Let P be a polytope in $\mathbb{R}^N$, and for each γ in P let ϕ_γ denote the map from $\mathbb{R}^N$ to $\mathbb{R}^N$ defined by*

$$\phi_\gamma(x) = x + (\gamma - v(x)), \tag{1.1}$$

‖ Current address: Department of Mathematics, University of Toronto, 100 St. George Street, Toronto, Ontario M5S 3G3, Canada.

where $v(x)$ is the closest vertex of P to x in Euclidean norm with some tie-breaking rule. Then for any compact set K there exists a bounded convex domain Q containing K and invariant in the sense that for all γ in P, $\phi_\gamma Q \subset Q$.

We encountered Theorem 1.1 in connection with a problem in digital color printing [**3**]. It turns out to have wide applicability (see, e.g., [**17, 23**]) and considerable mathematical depth.

The general problem that this theorem addresses concerns approximating a bounded one-sided sequence $\gamma(k)$ of arbitrary values in $\mathbb{R}^N$ by a sequence of elements $V(k)$ chosen from some finite set. In a wide range of applications, such as digital printing or scheduling, a reasonable sense of well-approximation is that there be a uniform bound (small if possible) on some norm of the *cumulative error vectors*

$$\varepsilon(n) = \sum_{k=0}^{n} (\gamma(k) - V(k)), \tag{1.2}$$

in which case the average error will go to zero. In applications it is assumed that the $V(k)$'s are taken as the vertices of a polytope P and the $\gamma(k)$'s belong to this polytope. In the case when the norm $\|\cdot\|$ is the sup norm and P the standard simplex, the problem of finding the best universal bounds on $\sup_n \|\varepsilon(n)\|$ is often referred to as the *chairman assignment problem*. This problem was posed by Niederreiter in [**22**] and considered in [**19, 20, 33, 34**]; we refer to the *generalized chairman assignment problem* when the norm is not necessarily the sup norm. Using the methods of [**17**] to find sequences that solve the generalized chairman assignment problem turns out not to be effective as it requires looking ahead with unbounded waiting time, at least for the sup norm. Waiting for future input data before making output decisions is not practical in many cases and sometimes even not acceptable as the scheduling algorithms may need not only to be *causal* (i.e. depend only on the past and the present and not also the future) but also to be *on the fly* (i.e. provide outputs essentially as soon as input data are known). Thus, if one is not looking for the best bound but only for a uniform one, or if one needs to perform continual optimization, then a surprisingly good approach is to write $\varepsilon(k)$ recursively as

$$\varepsilon(k+1) = \varepsilon(k) + \gamma(k+1) - v(\varepsilon(k) + \gamma(k+1)), \tag{1.3}$$

and use the *greedy algorithm* to specify $v(\varepsilon(k) + \gamma(k+1))$ which means here that $v(\varepsilon(k) + \gamma(k+1))$ is the vertex which minimizes the norm of $\varepsilon(k+1)$ with some tie-breaking rule: i.e. the vertex closest to $\varepsilon(k) + \gamma(k+1)$ with some tie-breaking rule. We shall be interested in the Euclidean norm; and in what follows, when we say *'the' greedy algorithm*, we mean the greedy algorithm using the Euclidean norm with any tie-breaking rule for points equidistant from two or more vertices. Later it will be convenient to express the greedy algorithm in terms of the notion of Voronoï regions. A *Voronoï region* of a vertex is the closure of the set of points that are closer to the vertex than to any other vertex (see Definition 2.1). Alternatively expressed, the greedy algorithm chooses the vertex in the Voronoï region containing $\varepsilon(k) + \gamma(k+1)$ with some decision rule when $\varepsilon(k) + \gamma(k+1)$ lies on a boundary.

Recursion (1.3) defines a non-autonomous dynamical system acting on the space of errors with a time-dependent parameter that belongs to the polytope. We are going to focus our attention on an associated dynamics in the ambient space of the polytope. We make the following changes of variables:

$$x(k) = \gamma(k) + \varepsilon(k-1), \tag{1.4}$$

and

$$V(k) = v(x(k)) = v(\gamma(k) + \varepsilon(k-1)). \tag{1.5}$$

Then, adding $\gamma(k+2)$ to both sides of (1.3) and reducing indices, we get

$$x(k+1) = x(k) + \gamma(k+1) - v(x(k)). \tag{1.6}$$

Recall that $v(x(k))$ is chosen as the closest vertex to $x(k)$ with some tie breaking rule. The orbit $x(k)$ can be expressed in terms of the mapping ϕ_γ in (1.1) by

$$x(k) = x(k-1) + (\gamma(k) - v(x(k-1))) = \phi_{\gamma(k)}(x(k-1)), \tag{1.7}$$

and bounding it in Euclidean norm is answered by Theorem 1.1. Bounding $\varepsilon(k)$ is equivalent to bounding $x(k)$, and an immediate corollary of Theorem 1.1 is the following convergence of averages:

$$\lim_{n\to\infty} \left\| \frac{1}{n}\sum_0^{n-1} \gamma(k) - \frac{1}{n}\sum_0^{n-1} v(x(k)) \right\| = 0. \tag{1.8}$$

1.1. *The greedy algorithm in digital printing.* To see the relevance of the above problem of approximating the $\gamma(k)$'s by the $V(k)$'s to color printing requires some background in color theory. The color printing alluded to concerns *digital halftoning* which is the printing technique for imitating full-color natural images (as opposed to colors placed at random on a page) by a checkerboard of dots limited to a few available colors. The simplest halftoning problem is that of obtaining gray-scale images with only black and white dots.

Color theory itself is a fascinating subject which we can only briefly touch upon. It has long been known that our perception of color has a vector space model. Electromagnetic spectra are so rich that to describe them requires an infinite-dimensional function space. Nevertheless, experiments indicate that the human visual system reduces the visible part of this space to a three-dimensional convex cone. In 1931 the *Commission Internationale de L'Eclairage* (CIE) devised a certain three-dimensional coordinate system called the *tristimulus space* which quantifies this cone. A complete account of how this coordinate system is contrived from color matching experiments is to be found in [**30**]. Our perception of light color is governed by a set of rules known as Grassman's laws, the consequence of which is that two different linear combinations of color vectors appear as the same color perhaps with different radiances, if their sums are collinear. Consider a section of paper on which a color image is to be printed using a digital printer, typically a laser or ink-jet printer which respectively use toner or ink to render colors. We shall side-step the non-mathematical practical difficulties involved in halftoning and assume an idealized situation: namely, the paper upon which an image is to be printed is partitioned into a checkerboard of tiny squares called *pixels* after a term used in color TV. We shall label pixel locations as we would the entries of a matrix. On each pixel is deposited a uniform color restricted to a very limited number of available choices which we shall call *pixel colors*. The problem is how should one choose pixel colors to render an acceptable full-color image.

Over a small area that still consists of a multitude of pixels the eye averages the received light. This means that we do not perceive colors on a page by taking positive linear combinations of color vectors, but rather convex combinations. The pixel colors commonly available to a printer are the standard inks or toner colors: cyan ($C = (21, 27, 72)$), magenta ($M = (33, 18, 22)$), yellow ($Y = (65, 76, 14)$). In addition, red ($R = (30, 18, 7)$), green ($G = (11, 22, 13)$), and blue ($B = (9, 7, 20)$) can be gotten by some kind of pairwise mixing of the above colors. Also available are the ink or toner for black ($K = (5, 6, 6)$), and white ($W = (84, 87, 105)$) which is taken for the color of paper in the absence of any ink or toner. In parentheses are typical tristimulus space coordinates for color vectors. The finite set $\{C, M, Y, R, G, B, W, K\}$ of ideally available printer color vectors are vertices of a convex polytope, in fact topologically a cube in three-space. For a general reference about how color theory affects digital color printing see [**13**].

Since color usually varies with low spacial frequency over many portions of natural images, it seems natural to simplify the halftoning problem by asking how one should choose pixel colors to represent a single uniform non-pixel color. Furthermore, since averaging plays such an essential role in color perception, it seems appropriate to recall the uniform distribution theory developed by T. van Aardenne-Ehrenfest, H. Davenport, K. F. Roth and others. Their results concerning the discrepancy between average values and expected ones can be applied to the problem of approximating arbitrary uniform grays over two-dimensional regions by combinations of black and white squares on a lattice. From a theorem of Roth (see [**6**, p. 3]) one can deduce that it is not possible to bound the cumulative error simultaneously for all gray levels and all rectangles. This is in contrast with printing in one dimension—i.e. printing on a line instead of a page—where it is well known, and as equation (1.10) below shows, that the cumulative error can be very well bounded on all segments, not only for gray levels but also for color. The discussion so far seems to suggest that, while there is no inherent difficulty in performing digital printing on a line, there may be one on a page, where it is of practical importance. However, the issue is not one of bounding cumulative errors, but rather of limiting their growth rate. Furthermore, a two-dimensional region R of pixels over which the eye averages is special.

For the sake of simplification, let R be an $n \times n$ square of pixels. The local cumulative error $E(R)$ made over the region R is given by

$$E(R) = \sum_{(i,j)\in R} (\gamma(i, j) - V(i, j)). \tag{1.9}$$

In the terminology of printing, the $\gamma(i, j)$'s are called *inputs* and the $V(i, j)$'s are called *outputs*. The object of halftoning is to choose outputs so that the average cumulative error $\|E(R)\|/n^2$ is small. With respect to a polytope P in color space whose vertices are the inputs, as we have seen, the greedy algorithm of equation (1.3) will achieve this. This is one of the methods of halftoning which is called *error diffusion* [**10, 11, 13, 27, 35**]. We call this particular version *simple error diffusion* [**3**]. Since this involves a sequential process, pixels must be linearly ordered, the usual order being the lexicographic one: namely, for an $M \times N$ array the (i, j)th pixel gets index $k = i \cdot N + j$ where the indexing starts at 0.

TABLE 1. Table of non-zero weights.

$w_m^k = p_{12}$	$w_{m-1}^k = p_{11}$	$w_{m-2}^k = p_{10}$	$w_{m-3}^k = p_9$	$w_{m-4}^k = p_8$
$w_{\ell-4}^k = p_7$	$w_{\ell-3}^k = p_6$	$w_{\ell-2}^k = p_5$	$w_{\ell-1}^k = p_4$	$w_\ell^k = p_3$
$w_2^k = p_2$	$w_1^k = p_1$	$(k+1)$st pixel	$\cdots \ell = k+3-N$	

If B is a bound on $\varepsilon(k)$, then $2B$ is a bound on the size of cumulative error

$$\varepsilon(L+n) - \varepsilon(L) = \sum_{j=L+1}^{L+n} (\gamma(j) - V(j)) \tag{1.10}$$

due to a run of n pixels starting at pixel L. Since there are n such runs in R, we have that $\|E(R)\| \leq 2Bn$, which leads to small average size of cumulative error for large enough n.

For two-dimensional digital printing simple error diffusion is deficient because, according to equation (1.3), even forgetting about the difficulty associated with changing lines, only pixel $k-1$ has a direct influence on pixel k. Other nearby pixels have no direct influence. Even using more exotic orderings cannot eliminate this flaw. However, one can simultaneously improve the influence of nearby pixels and deal with edge effects by using a *general error diffusion* which involves the greedy algorithm with respect to a weighting of past errors as follows. Let w_i^k, $1 \leq i \leq m$, be a system of weighting factors chosen in most cases for each k to be a probability vector. Indeed, we shall always assume that the weighting factors are chosen this way: i.e.

$$w_i^k \geq 0 \quad \text{and} \quad \sum_{i=1}^{m} w_i^k = 1.$$

Define

$$\varepsilon(k+1) = \left[\sum_{i=1}^{m} w_i^k \varepsilon(k+1-i)\right] + \gamma(k+1) - V(k+1) \tag{1.11}$$

where $V(k+1)$ minimizes $\|\varepsilon(k+1)\|$ with some tie-breaking rule. In the lexicographic ordering of pixels, even though those to the right and below have no influence on the current pixel, this method of weighting the past orbit seems to be a much better approximation to the eye's method of averaging than that of giving all the weight to the pixel immediately preceding the current one. A typical scheme for weighting factors for pixels not near borders is shown in Table 1 where p_i is a twelve-dimensional probability vector. Since this vector is allowed to depend on k, adjustments can be made to deal with pixels near borders.

The first such weighting scheme was introduced in [**10**] by R. Floyd and L. Steinberg in the seminal paper where they first described diffusion as a method for digital printing in 1975; this original scheme only involved p_1, p_5 and p_6, all other weighting factors being 0. Two more elaborate schemes can be found in [**11**] and [**31**] using the full Table 1, and a lot more have been devised and utilized since. In these two schemes, as is often the case now for square pixels, the highest weights are given to $p_1 = p_5$, the next highest to $p_4 = p_6$ and so on, with value decreasing with the distance to the kth pixel being treated (for specific values and other information concerning digital halftoning see the above references as well as [**13**] and [**35**]).

For general error diffusion the relation of $\varepsilon(k)$ to cumulative error is more complicated than equation (1.10). Nevertheless we shall show that, as in the case of simple error diffusion, for general error diffusion $\|E(R)\|/n^2 \to 0$.

First, we have the following corollary to Theorem 1.1.

COROLLARY 1.2. *For general error diffusion* $\|\varepsilon(k)\| < B$ *for some positive number* B.

Proof. Setting

$$x(k+1) = \left[\sum_{i=1}^{m} w_i^k \varepsilon(k+1-i)\right] + \gamma(k+1) \tag{1.12}$$

(in the literature this is called the *modified input*) we have

$$\varepsilon(k) = x(k) - v(x(k)),$$

where $v(x(k)) = V(k)$ is the closest vertex to $x(k)$ with some tie-breaking rule. Then

$$x(k+1) = \sum_{i=1}^{m} w_i^k [x(k+1-i) - v(x(k+1-i))] + \gamma(k+1) \tag{1.13}$$

$$= \sum_{i=1}^{m} w_i^k \phi_{\gamma(k+1)}(x(k+1-i)). \tag{1.14}$$

The fact that $x(k)$ of (1.13) is bounded, and equivalently that $\varepsilon(k)$ is bounded, is a consequence of the convexity of Q in Theorem 1.1. □

It is folklore knowledge in digital printing that for an $n \times n$ square R, the average error $\|E(R)\|/n^2$ goes to zero with n (see, e.g., [**8, 14**]). For the sake of completeness we indicate a proof: more precisely that $\|E(R)\| = O(n)$. Suppose that the pixel locations of the $n \times n$ square R contained in an $M \times N$ array, $N \gg n$, are $L + kK + j$, $0 \le j, k \le n-1$. Then from equation (1.11) we have

$$E(R) = \sum_{k=0}^{(n-1)} \sum_{j=L+kN}^{L+kN+n-1} (\gamma(j) - V(j))$$

$$= \sum_{i=1}^{m} w_i \sum_{k=0}^{n-1} \sum_{j=L+kN}^{L+kN+n-1} (\varepsilon(j) - \varepsilon(j-i)). \tag{1.15}$$

The only non-zero weights are $\{w_1, w_2, w_{pN+q} : 1 \le p \le 2, -2 \le q \le 2\}$.

Changing the limits on the index j, and interchanging order of summation (1.15), we get

$$E(R) = \sum_{i=1}^{m} w_i \sum_{j=0}^{n-1} \left[\sum_{k=0}^{n-1} \varepsilon(j+L+kN) - \sum_{k=0}^{n-1} \varepsilon(j+L+kN-i)\right].$$

Let

$$E_i \equiv \sum_{j=0}^{n-1} \left[\sum_{k=0}^{n-1} \varepsilon(j+L+kN) - \sum_{k=0}^{n-1} \varepsilon(j+L+kN-i)\right]. \tag{1.16}$$

Then for $0 \le p \le 2,\ -2 \le q \le 2$,

$$\begin{aligned} E_{pN+q} &= \sum_{j=0}^{n-1} \left[\sum_{k=0}^{n-1} \varepsilon(j+L+kN) - \sum_{k=0}^{n-1} \varepsilon(j+L+(k-p)N-q) \right] \\ &= \sum_{j=0}^{n-1} \sum_{k=0}^{n-1} \varepsilon(j+L+kN) - \sum_{j=-q}^{n-q-1} \sum_{k=-p}^{n-p-1} \varepsilon(j+L+kN). \end{aligned} \tag{1.17}$$

We leave to the reader the subsequent calculations, the result of which is

$$\|E(R)\| = O(n).$$

1.2. *A pursuit problem.* The following pursuit problem is another application of the main theorem. Since this section is not essential for the rest of the paper and the indexing can be tricky, readers might well skip it.

Suppose a predator is chasing its prey in a polytope, and the following conditions are imposed. The positions of the predator and prey at times $t_n = t_{n-1} + 1/n$ are denoted by $p(n)$ and $q(n)$ respectively (note that time intervals between movements are decreasing to zero yet sum to infinity). We take n to be greater than zero and assume that $p(0)$ and $q(0)$ are at two distinct corners, but any $p(0) \neq q(0)$ would do. Between times t_n and time t_{n+1} the prey moves in the direction of any point $\gamma(n+1)$ in the polytope with speed $n/(n+1)$ while the predator moves at unit speed but restricted towards some vertex, say $V(n+1)$, the choice of which being dictated by a pursuit strategy. Notice that the prey has the advantage over the predator of freer movement but the disadvantage of slower speed (although less and less so). The rules of movement are expressed by the following formulas:

$$p(n+1) = p(n) + \frac{1}{n}(V(n+1) - p(n)), \tag{1.18}$$

$$q(n+1) = q(n) + \frac{1}{n+1}(\gamma(n+1) - q(n)). \tag{1.19}$$

Let

$$\varepsilon(n) = n(q(n) - p(n+1)).$$

We have, using (1.19) at time t_{n+2} and (1.18) at time t_{n+1}, that

$$\varepsilon(n+1) = \varepsilon(n) + \gamma(n+1) - V(n+2). \tag{1.20}$$

As a capture strategy the predator chooses vertex $V(n+1)$ which minimizes $\varepsilon(n)$ with respect to the Euclidean norm, with some tie-breaking rule. In terms of the notation previously introduced,

$$V(n+1) = v(\varepsilon(n-1) + \gamma(n)) = v((n-1)(q(n-1) - p(n)) + \gamma(n)).$$

Thus,

$$\varepsilon(n+1) = \varepsilon(n) + \gamma(n+1) - v(\varepsilon(n) + \gamma(n+1)), \tag{1.21}$$

which we recognize as equation (1.3). It then follows from Theorem 1.1 that

$$\|q(n) - p(n+1)\| \to 0,$$

which implies

$$\|q(n) - p(n)\| \to 0.$$

In this sense the predator catches the prey.

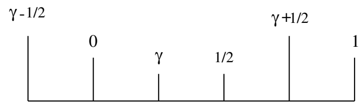

FIGURE 1. Diagram for simplest case.

1.3. *Instances of the greedy algorithm in classical mathematics.* The greedy algorithm with which we are concerned in this paper leads one into a realm of tremendous richness. Consider the simplest input case where $\gamma(k) \equiv \gamma$ is constant. The performance of the greedy algorithm reduces to the study of iterations of a single map ϕ_γ. This map is a piecewise isometry. In general the analysis of piecewise isometries can be of incredible depth and difficulty (see, e.g., [**4**]). Indeed, the simplest of all cases, the constant input case where the polytope $P = [0, 1]$ the unit interval, has significant mathematical substance. One easily checks that the interval $I = [\gamma - 1/2, \gamma + 1/2]$ is invariant and absorbing under ϕ_γ. Both the interval I and $[0, 1]$ are fundamental regions for the action of $\mathbb{Z}$ on $\mathbb{R}$ and

$$\phi_\gamma(x) = x + \gamma - v(x) = x + \gamma \mod \mathbb{Z}.$$

In other words, both I and $[0, 1]$ can be identified with the circle $\mathbb{R}/\mathbb{Z}$ and ϕ_γ with a rotation by an angle γ. This is one of the standard examples studied in ergodic theory and has the property of *unique ergodicity* when γ is irrational. In Figure 1 we have drawn the case in which $0 < \gamma < 1/2$. A consequence of unique ergodicity is that for irrational γ and any $x_0 \in I$, the ergodic averages of the characteristic function $\chi_{[1/2,\gamma+1/2]}(x_0 + k\gamma \mod 1)$ converges to the measure of $[1/2, \gamma + 1/2]$, which of course is equal to γ. For this particular interval, (1.8) and the equation

$$\chi_{[1/2,\gamma+1/2]}(x_0 + k\gamma \bmod 1) = \chi_{[1/2,\gamma+1/2]}(\phi_\gamma^k(x_0) = v(x(k)),$$

where $x(0) = x_0$, leads to the same result even for rational γ.

The type of sequence $v(x(k))$ of 0's and 1's generated by the recursion (1.6) is known as a *Sturmian sequence.* Such sequences were studied by Morse and Hedlund [**21**]. They have the property that there is a limiting frequency of 1's and the distribution of 0's and 1's is as even as possible. Arbitrary Sturmian sequences can be identified with orbits of homeomorphisms of the circle, but only those that appear for rigid rotations get generated by the greedy algorithm in the constant input case (see [**21**] and [**28**]). These sequences provide a symbolic counterpart to the theory of Diophantine approximation (see [**28**] and [**15**]) and have a long history. Johan Bernoulli was apparently the first to discover what came to be known as Sturmian sequences. He did that to describe an easier method of interpolation in astronomical tables [**7**]; Christoffel applied such sequences to modular arithmetic [**9**]; Smith used them to study well-distributed sequences of two symbols [**29**]. Both Christoffel and Smith seemed to have been unaware of the earlier work of Bernoulli. Bernoulli's work was later revisited by Markoff [**18**]. A generalization of Smith's view of the problem was considered in [**5**] where the circle was replaced by arbitrary-dimension tori, and the corresponding multidimensional continued fractions were described.

One more feature of the constant input case is the fact that under the action of ϕ_γ on $\mathbb{R}$ the interval I is invariant and *absorbing*—i.e. every orbit eventually enters and stays in I.

We expect that the following generalization holds in all dimensions, the proof of which we shall defer to a later work. For any lattice $\mathcal{L}$ with a preferred basis in $\mathbb{R}^N$, let P be the standard fundamental parallelepiped in the basis. Then for any $\gamma \in P$ there exists another fundamental region Q_γ which is an invariant absorbing set under the action of ϕ_γ; and on Q_γ the map ϕ_γ is a rotation by γ on $\mathbb{R}^N/\mathcal{L}$. The same results holds if P is a simplex instead of a parallelepiped, leading, however, to a different fundamental region Q_γ. Furthermore, there is a convergence of ergodic averages for special sets in the simplex case similar to that in the interval case. Finally, it is conceivable that this result could lead to higher-dimensional analogues to Sturmian sequences.

1.4. *Outline of the paper.* In §2 we describe how the proof of Theorem 1.1 evolved from dimension 2 to the general case. In §3 we prove Theorem 1.1 based on two propositions, Proposition 3.1 and Proposition 3.2, which we prove in §4. In §4, we show that the proof of Proposition 3.2 itself easily follows from a third proposition, Proposition 4.1. Proposition 4.1 posits the existence of a special subset of the $(N-1)$-dimensional sphere upon which the construction of the invariant set Q is based. We postpone the proof of that proposition to §5. Item (3) of that proposition is where all the difficulties and subtleties of our proof of Theorem 1.1 lie. Finally, at the end of §5 we have finished assembling all the elements for a complete proof of Theorem 1.1, and in §6 we present two Theorems concerning some additional properties which are byproducts of our method of proof.

2. *How the proof evolved from dimension two to the general case*

We shall adhere to the convention of writing the inner product of vectors as a dot product, the Euclidean norm as $\|\cdot\|$, and Euclidean distance between a pair of points or between a point and a set as $d(\cdot,\cdot)$.

Definition 2.1. Given a set of points $v_0,\dots,v_{M-1}$ in any dimension, the *Voronoï region* R_{v_i} of v_i is defined as

$$R_{v_i}=\bigcap_j \{x : \|x-v_i\| \le \|x-v_j\|\},$$

hence a polyhedron is the finite intersection of half-spaces. It is the set of points closer to v_i than to any other point in the set (or possibly no further than equidistant to some). Also let $v(x)$ denote *the closest vertex to x with a tie-breaking rule* such as choosing the vertex of smallest index in the case of ties.

Now let P be a convex polygon with vertices $v_0,\dots,v_{M-1}$. To make the definition of the map $\phi_\gamma(x)$ more specific, the tie-breaking rule for $v(x)$ will be for us the smallest index i such that $x \in R_{v_i}$. For example, if P is an interval $[v_0, v_1]$, then the Voronoï regions are

$$R_{v_0}=\left(-\infty,\frac{v_0+v_1}{2}\right]$$

and

$$R_{v_1}=\left(\frac{v_0+v_1}{2},\infty\right).$$

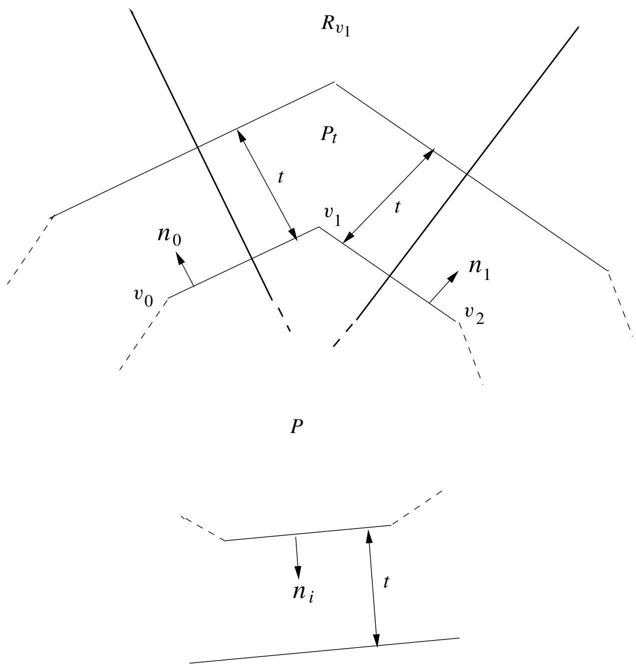

FIGURE 2. Polygons P, P_t, and Voronoï region R_{v_1}.

The mapping $\phi_\gamma : \mathbb{R} \hookrightarrow \mathbb{R}$ is given by

$$\phi_\gamma(x) = x + \gamma - v(x), \quad \text{where } v(x) = \begin{cases} v_0, & \text{if } x \le \dfrac{v_0 + v_1}{2}, \\ v_1, & \text{if } x > \dfrac{v_0 + v_1}{2}. \end{cases}$$

Theorem 2.1 and Theorem 2.2 are respectively dimension-one and -two (interval and polygon) versions of Theorem 1.1. Proofs of them appeared in [**3**].

THEOREM 2.1. *Let Q_t denote the interval generated by moving the ends of $P = [v_0, v_1]$ outward a distance $t \ge 0$: i.e. $Q_t = [v_0 - t, v_1 + t]$. Then for any $t \ge (v_1 - v_0)/2$ and $\gamma \in P$,*

$$\phi_\gamma(Q_t) \subset Q_t.$$

Let $v_i,\ i = 0, \dots, M-1,\ n \ge 2$ indexed in clockwise order P and $n_j,\ j = 1, \dots, n$ the unit normal vectors to edges $v_{j-1}v_j$ of P. Then $n_j \cdot (v_j - v_{j-1}) = 0$, and we can write $n_j \cdot v_{j-1} = n_j \cdot v_j = d_j$.

THEOREM 2.2. *Let Q_t denote the polygon generated by moving the edges of $P = \bigcap_j \{x : x \cdot n_j \le d_j\}$ perpendicularly outward a distance $t \ge 0$, i.e.*

$$Q_t = \bigcap_j \{x : x \cdot n_j \le d_j + t\}.$$

See Figure 2. There exists T such that for any $t \ge T$ and $\gamma \in P$

$$\phi_\gamma(Q_t) \subset Q_t.$$

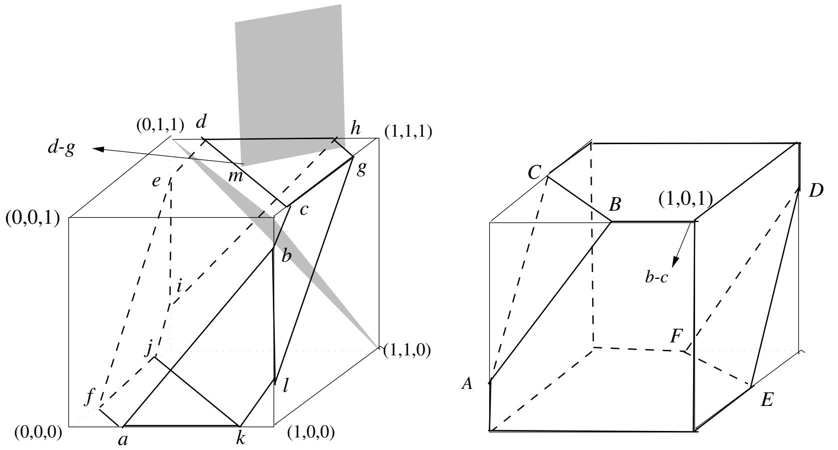

FIGURE 3. Polytope counterexample.

The above method of constructing an invariant Q by outwardly translating the planes of faces of a polytope P, even by different distances, is doomed to failure in dimension $N = 3$, as shown by the following counterexample.

Consider the octahedral polytope $abcdefghijkl$ in Figure 3. The faces $abcdef$ and $ghijkl$ are parallel to the plane $(0, 0, 0)(0, 1, 1)(1, 0, 1)$ cutting inside the cube as shown and equidistant from the points $(0, 0, 0)$ and $(1, 1, 1)$, respectively. The face of the Voronoï region R_g separating g from h is in the plane perpendicular to the top face of the cube through the midpoints of segments dc and hg. Invariance fails at the midpoint m of dc because here the vector d-g sticks out from the polytope. No matter how far the top plane of the cube is translated upward this vector will still stick out from the resulting polytope at such a point. So this means that the faces $abcdef$ and $ghijkl$ must be translated far enough apart. But then a new difficulty appears: the face of Voronoï region R_c lies in the plane determined by $(0, 1, 1)$, $(1, 0, 1)$ and the midpoint of bc. The vertex $(1, 0, 1)$ of the cube lies in this face, but here the vector b-c sticks out of this new candidate for an invariant polytope. On one hand this second difficulty cannot be overcome by moving the front and side faces outward, but on the other hand moving the top face up reintroduces the first difficulty. We leave it to the reader to verify that one problem cannot be overcome without introducing the other: for details and the general case of dimension $N > 2$, see [**25**].

So another idea for a general construction of an invariant set is needed (special constructions for special polytopes are discussed in [**16**]). Imagine the limiting case of a sphere centered at the origin scaled down to the unit sphere as the radius goes to infinity. The edges at a vertex are normals to hyperplanes through the origin. These hyperplanes bound the limiting Voronoï region at that vertex. The following, which we shall verify later for all dimensions, is easy to see in the two-dimensional case below. It is the most important element of the proof. *The cone of normals to the supporting hyperplanes at a vertex agrees with the Voronoï region at infinity of that vertex* (see the proof of Proposition 3.1). Thus, the tangent space of the sphere at a normal in the Voronoï region of a vertex is a supporting

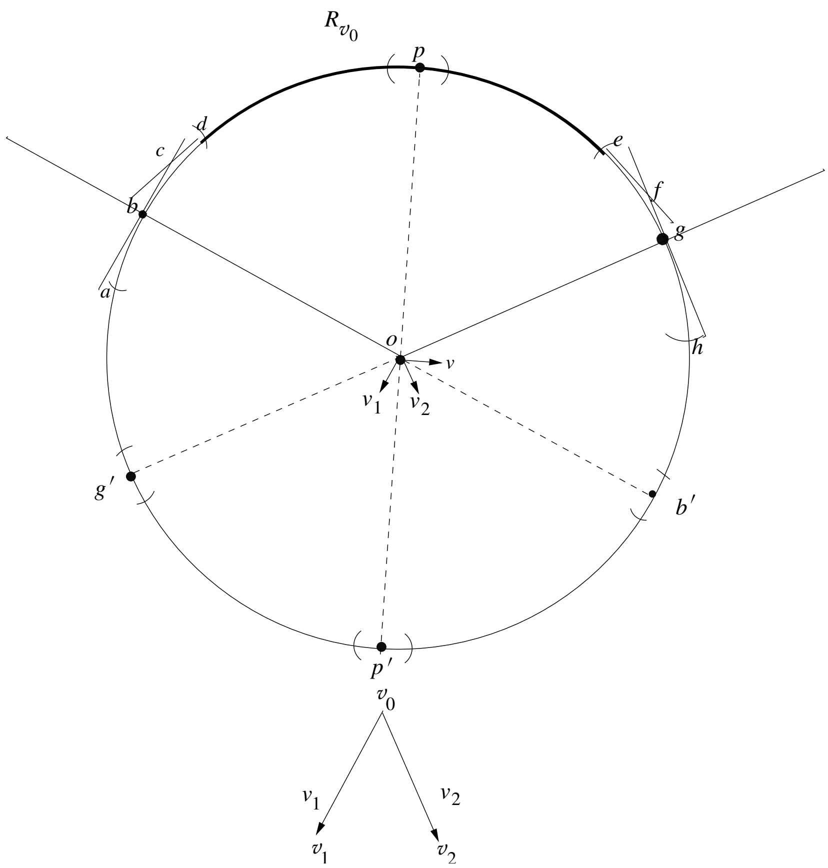

FIGURE 4. Construction of Q in dimension two.

hyperplane of the polytope at that vertex. In this idealized situation the polytope has shrunk to a point at the origin; and the unit sphere (equivalently the sphere at infinity) is invariant. Backing off from infinity to a sufficiently large sphere, the invariance fails only on neighborhoods of Voronoï boundaries. Resolving this difficulty then becomes the main task of the proof. To do this we construct a special subset Ω of the unit sphere the tangent planes of which are supporting hyperplanes of a *proto-invariant* convex set Q_∞. Then the set $Q = \rho Q_\infty$ for large enough ρ will be the desired invariant convex set.

Before launching into the formal proof for all dimensions, we illustrate the idea for some low-dimensional polytopes indicating the induction upon which the general construction is based.

First, we take up a construction of Q via a set Ω in dimension two. Consider the idealized limiting circle scaled down to the unit circle as depicted in Figure 4. Let v_1, v_0, v_2 be three successive vertices of a polygon, and $\nu_1 = v_1 - v_0$ and $\nu_2 = v_2 - v_0$. The points b and g lie the intersection of rays perpendicular to ν_1 and ν_2 with S. These rays are parallel to the unbounded edges of the Voronoï region R_{v_0} for the vertex v_0.

Using the same letters to denote corresponding points on a dilated circle ρS, the intersection of the infinite bounding edges of the actual Voronoï region with the circumference of the dilated circle for large ρ would be in the arc ad and in the arc eh. The only place a vector $\gamma - v_0$ would stick out of the circle for some $\gamma \in P$ would be at points b and g (e.g. $\gamma - v_1$ at b) and nearby points relatively speaking. So, in order to avoid this difficulty, we remove small open arcs ad and eh of S, retaining the points b and g, and carry out this procedure for each Voronoï region R_{v_j}. The size of the arcs is subject to the restriction that the removed arcs do not interfere with one another. It can then be seen that, if Q_∞ is the convex set supported by tangent lines to the unit sphere at points of Ω, then the convex set ρQ_∞ is invariant for large enough ρ.

A construction for Q via a set Ω in dimension 3 is similar. For the sake of simplicity, suppose that three edges meet at vertex v_0 and that v_1, v_2 and v_3 are the neighboring vertices, as in the bottom of Figure 5. Let $\nu_1 = v_1 - v_0$, $\nu_2 = v_2 - v_0$ and $\nu_3 = v_3 - v_0$ be the normals to the faces of the Voronoï regions R_{v_0}. The planes defined by these normals intersect the idealized scaled limiting unit sphere in the spherical triangle abc. To construct Ω we remove from the unit sphere small open spherical caps centered at a, b and c, but retain these centers. In the complement of what has been removed, we remove open strips centered on the edges of the spherical triangle, but retain their center-lines which in this case are the closed arcs de, fg and hk. In Figure 5 we have illustrated the construction for Voronoï region R_{v_0}. We carry out this procedure for each Voronoï region. Again, caps and strips should be chosen so that they do not interfere with one another. The set Ω is what survives. As before, it can be seen that the failure of invariance of large spheres is overcome by the set Ω, and the convex set Q_∞ supported by tangent planes to the unit sphere at points of Ω, whereupon $Q = rQ_\infty$ is invariant for large enough r.

Note that the boundary of the removed caps from the 2-sphere in three dimensions has the same structure as Ω in two dimensions. Also, the boundary of a cross-section of a removed strip from the sphere has the same structure as a line through a boundary of a removed arc on the circle (like the situation on the arc ad). Thus, we get a glimpse of the induction about to be carried out.

In dimension four, the intersection of a Voronoï region and the 3-sphere is a spherical polytope. To construct Ω we proceed exactly as before. We remove small open 4-spheres about the vertices of the polytope retaining the centers. From what remains, we remove open tubes around the edges, retaining the edges. Finally, from what still remains of the unit sphere we remove slabs about the faces of the polytope while retaining the faces. We do this for each Voronoï region. Again, a restriction on the removed sets is that they do not interfere with one another.

Again, we note that the boundary of a removed spherical neighborhood from the 3-sphere in four dimensions is a 2-sphere having the same structure as the three-dimensional Ω. The retained center of such a removed spherical neighborhood is not visible in the three-dimensional subspace in which the bounding 2-sphere lives. Rather, the center of this 2-sphere is a projection of the retained point in Ω. This corresponds to what is happening in the previous dimension. In Figure 5 the point a is not in the plane which contains the boundary circle of the removed spherical cap on the 2-sphere. The center of this circle is the projection of the point a to this plane. Going down one more dimension,

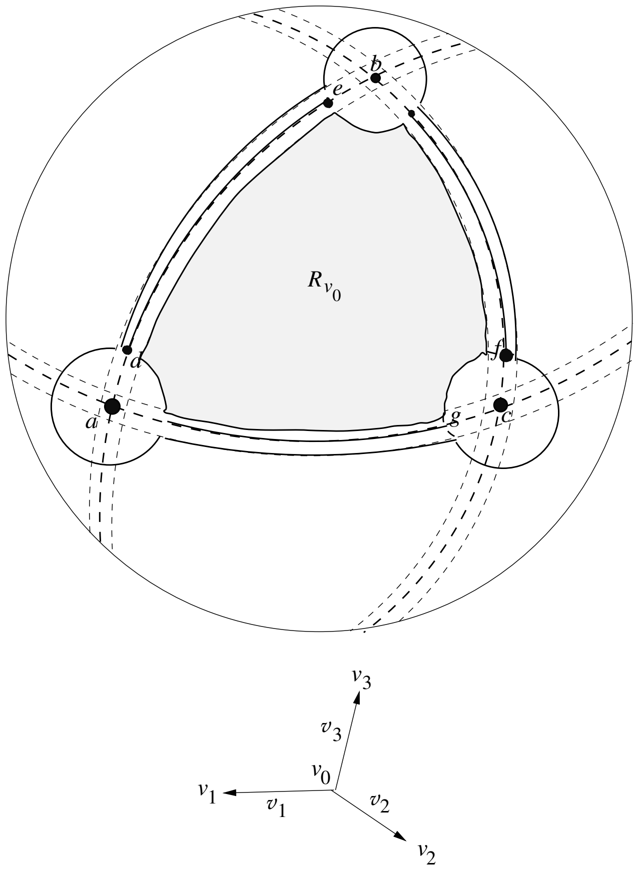

FIGURE 5. Construction of Q in dimension three.

this corresponds to the fact that b does not lie on the line through a and d. A cross-section of a removed tube has the same structure as the boundary of a removed cap in dimension three which in turn has the same structure as the removed circle in dimension two.

Remark. Returning to Figure 4 we note that more points could have been removed between d and e without compromising invariance. Similarly, the same can be said for removing more points from Ω in any higher dimension. Our actual construction of Ω will take advantage of this fact, and we shall remove much more from the unit sphere than what we have just described. This has two virtues. First, describing removed sets is slightly easier. Second, it leads to a simplified proof from which we get for free Theorem 6.2.

For the construction in dimension 2 of this more truncated Ω, let p and p' be the points of intersection of the unit sphere with the line normal to $\nu = v_i - v_j$ for a pair of vertices v_i and v_j. Remove from the unit sphere small open arcs (symmetric) about p and p' but retaining p and p'. Do this for *every* pair of vertices v_i and v_j of the polytope. The only restriction is that removed arcs be small enough not to interfere with each other.

The set Ω is smaller than in the previous construction which results in fewer supporting lines for Q_∞ which then results in a larger Q_∞. Nevertheless, the ultimate Q will still be invariant.

To construct Ω in dimension 3, consider all possible intersections with the unit sphere of planes whose normals are given by $v_i - v_j$ for every distinct pair of adjacent vertices $v_i \neq v_j$ of the polytope P. These intersections consist of one-dimensional lines, and two-dimensional planes. We first remove caps centered at all the points that are intersections of the one-dimensional lines with the unit sphere while retaining the center-points. From what is left we remove all strips centered at great circle arcs that are intersections of two-dimensional planes with the unit sphere while retaining the central arcs. Once again the removed sets must be at least small enough so that they do not interfere with one another. To facilitate our proof it turns out to be convenient to make these removed sets even smaller than just the requirements for non-interference. As before, the resulting Ω leads to an invariant Q.

The above will be the procedure for constructing Ω in all dimensions. Showing that a supporting hyperplane tangent to the unit sphere at a point in Ω is not cut off at the wrong place by another one is delicate. This is clear in dimension 2, but far from obvious in higher dimensions.

3. *Proof of Theorem 1.1*

We shall use freely some notations, definitions, and basic results from classical convex geometry (see, for instance, [**26**] or [**27**] and references therein).

3.1. *Notation.*

- P is a polytope namely, the convex hull of a finite set (vertices) $\{v_1, \ldots, v_M\}$ of extreme points (vertices) in $\mathbb{R}^N$. By the classical *finite basis theorem for polytopes* attributed to Minkowski (1896), Steinitz (1916) and Weyl (1935) (see [**27**, p. 89]), P is a bounded polyhedron.
- $C(v_i)$ denotes the cone of outward normals to the supporting hyperplanes of P at vertex v_i; and $C_1(v_i)$ the unit normals in $C(v_i)$. See spherical triangle abc in Figure 5.
- As a polyhedron, P can be expressed by

$$P = \bigcap_{\text{all } v_i} \bigcap_{\omega \in C(v_i)} \{x : \omega \cdot (x - v_i) \leq 0\}. \tag{3.1}$$

3.2. *Proof of the main result.*

We shall assume two propositions, the proofs of which shall be deferred to the next two sections.

PROPOSITION 3.1. *For each $\varepsilon > 0$ there exists $D > 0$ such that if $\|x\| > D$ then $d(x/\|x\|, C_1(v(x))) < \varepsilon$.*

PROPOSITION 3.2. *There exists a closed subset Ω of the unit sphere in $\mathbb{R}^N$ centered at the origin with the following properties:*

(1) *Ω is not confined to any half-sphere;*

(2) *there exists an* ε, $0 < \varepsilon < 1$, *for all* $i = 1, \ldots, M$ *such that if* $y \in C_1(v_i)$ *and* $\omega \in \Omega \setminus C_1(v_i)$ *then there exists* $\omega' \in C_1(v_i) \cap \Omega$ *such that*

$$d(\omega', y) < d(\omega, y) - \varepsilon. \tag{3.2}$$

Proof of Theorem 1.1. Let

$$Q_\infty \equiv \bigcap_{\omega\in\Omega} \{x : x \cdot \omega \leq 1\}, \tag{3.3}$$

and

$$\rho Q_\infty = \bigcap_{\omega\in\Omega} \{x : x \cdot \omega \leq \rho\}. \tag{3.4}$$

We shall show that for all sufficiently large ρ the set ρQ_∞ satisfies Theorem 1.1. It is clear by definition that Q_∞ and hence ρQ_∞ are convex. If Q_∞ were not bounded, it would contain a ray $\mathbb{R}^+ x = \{tx : t \geq 0\}$, and the half-sphere determined by the hyperplane normal to x would contain Ω, contradicting Proposition 3.2. Thus Q_∞ and ρQ_∞ are bounded.

We must show that there exists an $R > 0$ such that

$$\omega \cdot \phi_\gamma(x) = \omega \cdot (x + \gamma - v(x)) \leq \rho \tag{3.5}$$

for all $x \in \rho Q_\infty$, $\rho \geq R$, $\gamma \in P$ and $\omega \in \Omega$.

Case I: $\omega \in C(v(x))$. By (3.1) and (3.4), for $\gamma \in P$ we have

$$\omega \cdot x + \omega \cdot (\gamma - v(x)) \leq \rho. \tag{3.6}$$

The inequality is satisfied no matter what the size of $\|x\|$ and $\rho > 0$.

Case II: $\omega \notin C(v(x))$. Let ε be the one given in Proposition 3.2. From Proposition 3.1 there exists $R_0 > 0$ such that for $\|x\| > R_0$ there exists $y \in C_1(v(x))$ satisfying

$$d\left(\frac{x}{\|x\|}, y\right) < \varepsilon/4.$$

By Proposition 3.2 there exists $\omega' \in C_1(v(x)) \cap \Omega$ such that

$$d(w, y) > d(\omega', y) + \varepsilon.$$

By the triangle inequality

$$\begin{aligned} d\left(\omega, \frac{x}{\|x\|}\right) &\geq d(\omega, y) - d\left(y, \frac{x}{\|x\|}\right) \\ &> d(\omega', y) + \varepsilon - d\left(y, \frac{x}{\|x\|}\right) \\ &\geq d\left(\omega', \frac{x}{\|x\|}\right) + \varepsilon - 2d\left(y, \frac{x}{\|x\|}\right) \\ &\geq d\left(\frac{x}{\|x\|}, \omega'\right) + \varepsilon/2. \end{aligned}$$

Let θ be the angle between x and ω, and θ' between x and ω'. Using the fact that differences in arc lengths are greater than differences in corresponding chord lengths (an exercise hint: make chords parallel), we get

$$\theta - \theta' > d\left(\omega, \frac{x}{\|x\|}\right) - d\left(\frac{x}{\|x\|}, \omega'\right) > \varepsilon/2 \tag{3.7}$$

for the arc-length difference. Thus,

$$\cos(\theta) < \cos(\theta' + \varepsilon/2) < \cos(\theta')\cos(\varepsilon/2), \tag{3.8}$$

and we arrive at

$$\omega \cdot \frac{x}{\|x\|} < \omega' \cdot \frac{x}{\|x\|}(1 - e),$$

where $e = 1 - \cos(\varepsilon/2)$. Note that $0 < e < 1$, depending only on ε. We then get

$$\omega \cdot x + \omega \cdot (\gamma - v(x)) < \omega' \cdot x(1 - e) + \mathrm{diam}(P). \tag{3.9}$$

We can therefore choose $R_1 > R_0$ such that for $\|x\| \geq R_1$ the right-hand side of (3.9) is $< \rho$. Let

$$R = R_1 + \mathrm{diam}(P), \tag{3.10}$$

and let $x \in \rho Q_\infty$ where $\rho > R$. By what we have just shown, if $\|x\| \geq R_1$ then x satisfies (3.5): so $\phi_\gamma(x) \in \rho Q_\infty$. Finally, if $\|x\| < R_1$ then $x \in \rho Q_\infty$, because $\|x - \phi_\gamma(x)\| \leq \mathrm{diam}(P)$. □

Remark 3.1. At this point we have also proved that given any compact set K the set Q can be chosen to contain K.

4. *Proofs of Propositions 3.1 and 3.2*

4.1. *Proof of Proposition 3.1.* By a standard argument in linear programming—namely, for each vector u the function $u \cdot x$ achieves a maximum over P at some vertex v_i which means that $u \cdot (x - v_i) \leq 0$ for $x \in P$, or equivalently that $u \in C(v_i)$—we have $\mathbb{R}^N = \bigcup_i C(v_i)$.

Furthermore, by the *decomposition theorem of polyhedra* due to Motzkin (1936) (see [**27**, p. 88]) a Voronoï region R_{v_i}, being a polyhedron, can be written

$$R_{v_i} = P_{v_i} + K_{v_i},$$

where P_{v_i} is a polytope (recall a polytope is a bounded set) and K_{v_i} is a polyhedral cone. Since $\mathbb{R}^N = \bigcup_i R_{v_i}$, it is easy to prove that $\mathbb{R}^N = \bigcup_i K_{v_i}$ as well.

We show that $K_{v_i} = C(v_i)$ as follows. Let $u \in C(v_i)$, and $t \geq 0$. The vector u determines a supporting hyperplane for P at v_i which implies that v_i is the closest vector of P to the vector $tu + v_i$: so $tu + v_i \in R_{v_i}$. Thus, from the polyhedral decomposition theorem $tu + v_i = p(t) + k(t)$ for all $t \geq 0$ where $p(t) \in P_{v_i}$ and $k(t) \in K_{v_i}$. Therefore,

$$u = \frac{p(t) - v_i}{t} + \frac{k(t)}{t}. \tag{4.1}$$

The first term on the left in (4.1) converges to 0 as $t \to \infty$ so that the second term, which always belongs to $K(v_i)$ by convexity, is bounded for all $t > 0$. Thus, there is a convergent subsequence whose limit is in $K(v_i)$, which implies that $u \in K_{v_i}$. Hence $C(v_i) \subset K_{v_i}$.

Since the two collections of $C(v_i)$'s and of K_{v_i}'s form partitions of $\mathbb{R}^N$, the reverse inclusion follows easily from the fact that the interiors of $C(v_i)$ and $C(v_j)$ are disjoint for $i \neq j$, and similarly for the interiors of K_{v_i} and K_{v_j}.

Given any $x \in \mathbb{R}^N$, we have just shown

$$x \in P_{v(x)} + C(v(x)),$$

so there exists $p(x) \in P_{v(x)}$ such that $x - p(x) \in C(v(x))$. Theorem 3.1 now follows from the easily proved limit

$$\left\| \frac{x}{\|x\|} - \frac{x - p(x)}{\|x - p(x)\|} \right\| \to 0, \quad \text{as } \|x\| \to \infty. \qquad \square$$

4.2. *Proof of Proposition 3.2.* The proof of Proposition 3.2 follows easily from the following proposition, the proof of which has been postponed until the last section. As previously stated, this proposition is the heart of the matter.

PROPOSITION 4.1. *There exist a partition of each $C_1(v_i)$ into a finite number of connected sets σ and a closed subset Ω of the unit sphere S in $\mathbb{R}^N$ centered at the origin with the following properties:*

(1) *Ω is not confined to any half-sphere;*

(2) *for each $\sigma \subset C_1(v_i)$ there is a $\varepsilon_\sigma > 0$ such that*

$$\Omega \cap U_{\varepsilon_\sigma}(\bar{\sigma}) \subset \bar{\sigma},$$

where $U_\varepsilon(A) = \{x \mid d(x, A) < \varepsilon\}$ denotes the ε-neighborhood of a set A;

(3) *the closest points $\omega \in \Omega$ to $y \in \bar{\sigma} \subset C_1(v_i)$ are in the same set closure $\bar{\sigma}$. In particular, $\bar{\sigma} \cap \Omega \neq \emptyset$.*

The sets $\sigma \in C_1(v_i)$ will be chosen as *cells* of a simplicial decomposition induced by great spheres on S as specified in §5.

Proof of Proposition 3.2. We must prove that there exists ε such that for all v_i and all $y \in C_1(v_i)$

$$d(y, \Omega \setminus C_1(v_i)) - d(y, C_1(v_i) \cap \Omega) > \varepsilon. \tag{4.2}$$

For $\sigma \subset C_1(v_i)$ define a function $f_\sigma : \bar{\sigma} \to \mathbb{R}$ by

$$f_\sigma(y) \equiv d(y, \Omega \setminus U_{\varepsilon_\sigma}(\bar{\sigma})) - d(y, C_1(v_i) \cap \Omega),$$

where ε_σ is given by Proposition 4.1. Then f_σ is continuous and strictly positive on a compact set: hence assumes a minimum $e_\sigma > 0$. We claim $\varepsilon = \min e_\sigma$, where the minimum is taken over all $\sigma \subset C_1(v_i)$ and all v_i's, satisfies (4.2). Observe that

$$d(y, \Omega \setminus C_1(v_i)) - d(y, C_1(v_i) \cap \Omega) \geq d(y, \Omega \setminus \bar{\sigma}) - d(y, C_1(v_i) \cap \Omega), \tag{4.3}$$

where $y \in \bar{\sigma} \subset C_1(v_i)$. Now by (2) and (3) of Proposition 4.1 the right-hand side of (4.3) is equal to $d(y, \Omega \setminus U_{\varepsilon_\sigma}(\bar{\sigma})) - d(y, C_1(v_i) \cap \Omega) = f_\sigma(y) \geq \varepsilon$. $\square$

5. *Proof of Proposition 4.1*

5.1. *Notation.*

- $S = S(c, r) \subset \mathbb{R}^{n+1} = \mathbb{R}^N$ denotes an n-dimensional sphere with radius r and center c. We shall call the intersection of S with any k-dimensional affine subspace of $\mathbb{R}^{n+1}$ containing c a *great sphere of dimension* $k-1$ or simply a *great sphere.*
- For A a closed subset of S, the set $\mathrm{clos}(x, A)$ of points in A closest to a point $x \in S$ is given by
$$\mathrm{clos}(x, A) = \{a \in A : \forall b \in A,\ |a - x| \leq |b - x|\}.$$
- $\mathcal{V} = \{V_1, V_2, \dots, V_s\}$ denotes a collection of codimension-one affine subspaces of $\mathbb{R}^N$ containing the center of S. For any x in S, $\mathcal{V}_x$ will stand for the collection of V_i's that contain x.
- For each α a subset of $\{1, 2, \dots, s\}$ we define the subspace
$$V_\alpha = \bigcap_{i \in \alpha} V_i$$
and the great sphere in S
$$S_\alpha = S \cap V_\alpha.$$
By the usual convention, $V_\emptyset = \mathbb{R}^N$ and $S_\emptyset = S$.
- We call the pair $\mathbb{S} = (S, \mathcal{V})$ a *stratification* of S induced by the subsets of $\mathcal{V}$. We call the lowest-dimensional S_α containing a point x the *stratum of* x and denote it by $S_\alpha(x)$. It is well defined, since the intersection of any such two strata would be a lower-dimensional one. A stratification induces a partition of S into *cells* of dimensions $0, 1, \dots, \dim S$, the *cell containing* x, denoted $\sigma(x)$, being the connected component of $\{y \in S_\alpha(x) \mid S_\alpha(y) = S_\alpha(x)\}$ containing x. Two stratifications $\mathbb{S}_1 = (S_1, \mathcal{V}_1)$ and $\mathbb{S}_2 = (S_2, \mathcal{V}_2)$ are called *isometric* if there exists a bijection $j : \mathcal{V}_1 \to \mathcal{V}_2$ and an orthogonal affine map $O : S_1 \to S_2$ such that
$$O(S_1 \cap V_i) = S_2 \cap j(V_i)$$
for every $V_i \in \mathcal{V}_1$.
- Let $V_\alpha^\perp \subset \mathbb{R}^N$ be the affine subspace orthogonal to $V_\alpha \subset \mathbb{R}^N$ passing through the center of S and $V_\alpha^\perp(\varepsilon) \equiv U_\varepsilon(c) \cap V_\alpha^\perp$.
- The *ε-tubular neighborhood of S_α* is defined as
$$T_\varepsilon(S_\alpha) = (V_\alpha \times V_\alpha^\perp(\varepsilon)) \cap S. \tag{5.1}$$
When $\dim S_\alpha = k$ we also call $T_\varepsilon(S_\alpha)$ a *k-tube*.
- For every $z \in S_\alpha$ we denote the great sphere orthogonal to S_α at z by $S_\alpha^\perp(z)$. Notice that
$$S_\alpha^\perp(z) = (\mathbb{R} \cdot (z - c) \times V_\alpha^\perp) \cap S.$$
- For every $z \in S_\alpha$ it is easy to check that the set
$$F_\alpha(z, \varepsilon) = (\mathbb{R}^+ \cdot (z - c) \times V_\alpha^\perp(\varepsilon)) \cap S.$$
is a spherical cap. We call it an *(α, ε)-spherical cap with top at z* or simply a *spherical cap*.

- For every $z \in S_\alpha$ we define a neighborhood fibered by (α, ε)-spherical caps

$$P_\alpha(z, \varepsilon) = \bigcup_{y \in U_\varepsilon(z) \cap \sigma(z)} F_\alpha(y, \varepsilon), \tag{5.2}$$

  called the (α, ε)*-patch centered at* z where α is given by the stratum containing z.
- The boundary $\partial_{\mathrm{rel}} F_\alpha(z, \varepsilon)$ of the spherical cap $F_\alpha(z, \varepsilon)$ relative to the subspace $S_\alpha^\perp(z)$ is given by

$$\partial_{\mathrm{rel}} F_\alpha(z, \varepsilon) = (\mathbb{R}^+ \cdot (z - c) \times (\partial U_\varepsilon(c) \cap V_\alpha^\perp)) \cap S.$$

  The relative boundaries of spherical caps are Euclidean spheres though not great ones: rather they are analogous to meridians with the north pole at z. Note that

$$\dim \partial_{\mathrm{rel}} F_\alpha(z, \varepsilon) = n - \dim S_\alpha - 1 = \dim S - \dim V_\alpha. \tag{5.3}$$

- The stratification $\mathbb{S}$ induces the stratification $(\partial_{\mathrm{rel}} F_\alpha(z, \varepsilon), \mathcal{V}_z)$ on the sphere $\partial_{\mathrm{rel}} F_\alpha(z, \varepsilon)$.

We leave to the reader to check the following relations:

$$T_\varepsilon(S_\alpha) = \bigcup_{z \in S_\alpha} F_\alpha(z, \varepsilon), \tag{5.4}$$

otherwise the (α, ε)-spherical caps foliate the ε-tubular neighborhoods of S_α;

$$F_\alpha(z, \varepsilon) \subset U_{(\sqrt{2})\varepsilon}(z); \tag{5.5}$$

$$P_\alpha(x, \varepsilon) \subset U_{(1+\sqrt{2})\varepsilon}(x); \tag{5.6}$$

$$S = \bigcup_{z \in S_\alpha} S_\alpha^\perp(z); \tag{5.7}$$

and

$$\partial_{\mathrm{rel}} F_\alpha(z, \varepsilon) = S\left(c + \frac{\sqrt{r^2 - \varepsilon^2}}{r}(z - c), \varepsilon\right) \cap S_\alpha^\perp(z).$$

We also leave as an exercise the following.

LEMMA 5.1. (Spherical cap characterization) *A spherical cap is the intersection of a half-space and a sphere, and conversely.*

5.2. *Spherical caps, the set Ω and the Main Lemma.*

LEMMA 5.2. (Spherical cap boundary isometry) *If $\mathcal{V}_x = \mathcal{V}_y$ and $x, y \in S_\alpha(x)$, then $(\partial_{\mathrm{rel}} F_\alpha(x, \varepsilon), \mathcal{V}_x)$ and $(\partial_{\mathrm{rel}} F_\alpha(y, \varepsilon), \mathcal{V}_y)$ are isometric stratifications for any $\varepsilon > 0$.*

Proof. Consider the plane of x, y and c. This plane is a subspace of V_α. Then, the affine orthogonal map O, which rotates this plane about c, rotating y to x, and which is the identity on the space orthogonal to the plane, is the required isometry. □

LEMMA 5.3. (Spherical cap intersection) *Let $S_{\alpha'} \subset S_\alpha$. Given $\mu < r$, if $x \in S_\alpha$, and $z \in S_{\alpha'}$, then either*

$$S_\alpha^\perp(x) \cap F_{\alpha'}(z, \mu) = \emptyset,$$

or

$$S_\alpha^\perp(x) \cap F_{\alpha'}(z,\mu) = F_\alpha(x,\delta), \tag{5.8}$$

where δ and μ are related by

$$r^2 - \mu^2 = (r^2 - \delta^2)\left(\frac{(x-c)\cdot(z-c)}{r^2}\right)^2, \tag{5.9}$$

with r standing for the radius of the sphere S centered at c.

Moreover, if $\delta/\mu \le \sqrt{2\sqrt{2}-2} \simeq 0.911$ and $\delta/\mu \le \tau$, then

$$d(x, \partial_{\mathrm{rel}} F_{\alpha'}(z,\mu)) < \tau\delta. \tag{5.10}$$

Proof. Observe that if $x \notin F_\alpha(x,\mu)$ then $S_\alpha^\perp(x) \cap F_{\alpha'}(z,\mu) = \emptyset$. Now let $x \in F_\alpha(x,\mu) \setminus \partial_{\mathrm{rel}} F_{\alpha'}(z,\mu)$ and $x \ne z$. Then $S_\alpha^\perp(x) \cap F_{\alpha'}(z,\mu) \ne \emptyset$. By Lemma 5.1 $S_\alpha^\perp(x) \cap F_{\alpha'}(z,\mu)$ is a spherical cap and we are justified in denoting it by $F_\alpha(x,\delta)$. We must now determine δ.

Choose $v \in \partial F_{\alpha'}(z,\mu)$. Then $v \in \partial_{\mathrm{rel}} F_\alpha(x,\delta) \cap \partial_{\mathrm{rel}} F_{\alpha'}(z,\mu)$. Then the vectors $v-c,\ x-c,\ z-c$ determine the three-dimensional c-centered sphere S_3 that contains the points v, x and z. We assume that c is at the origin and the point z is at the north pole. The intersection of $F_{\alpha'}(z,\mu)$ with S_3 is a spherical cap whose boundary is a meridian which is also $\partial_{\mathrm{rel}} F_{\alpha'}(z,\mu) \cap S_3$; this meridian is determined by μ and the intersection of $\partial_{\mathrm{rel}} F_\alpha(x,\delta)$ with S_3. This intersection is an arc of a great circle whose endpoints lie on this meridian. Let y be the intersection of the great circle determined by x and z with this meridian. The projections of v on x and z are $p = (v\cdot x/r^2)x$ and $q = (v \cdot z/r^2)z$, respectively. Note that

$$\frac{v\cdot z}{r^2} = \sqrt{1 - \left(\frac{\mu}{r}\right)^2} \tag{5.11}$$

and that

$$\frac{v\cdot x}{r^2} = \sqrt{1 - \left(\frac{\delta}{r}\right)^2}. \tag{5.12}$$

We can write

$$v = \left(\frac{v\cdot x}{r^2}\right)x + u, \tag{5.13}$$

which exhibits the contribution $u \in V_\alpha^\perp \subset V_{\alpha'}^\perp$ to v and yields

$$v\cdot z = \left(\frac{v\cdot x}{r^2}\right)x\cdot z. \tag{5.14}$$

Equation (5.9) then follows by substituting (5.11) and (5.12) in (5.14).

For the case $x = z$ it is apparent that (5.9) holds with $\delta = \mu$.

Finally, for the case $x \in \partial_{\mathrm{rel}} F_{\alpha'}(z,\mu)$ it is easy to see that (5.9) holds with $\delta = 0$.

To obtain equation (5.10) we first prove that y is the closest vector to x in $\partial_{\mathrm{rel}} F_{\alpha'}(z,\mu)$. Let v play the role of any vector in $\partial_{\mathrm{rel}} F_{\alpha'}(z,\mu)$, not just one of those at a distance δ from the axis of S through x, and observe that $d(x,v) > d(x,y)$.

From the fact that angle xpy is acute—i.e.

$$d(x,y)^2 \le d(y,p)^2 + d(p,x)^2, \tag{5.15}$$

and the fact that the shortest distance to the circumference is along a radius—i.e.

$$d(p,x) \leq d(p,y), \tag{5.16}$$

we get

$$d(x,y) \leq \sqrt{2} d(y,p). \tag{5.17}$$

Furthermore, from

$$d(y,p) = d(y,q) - d(p,q),$$

and the fact that $\triangle vpq$ is a right triangle we have that

$$\begin{aligned} d(x,y) &\leq \sqrt{2}(\mu - \sqrt{\mu^2 - \delta^2}) \\ &= \sqrt{2}\mu \left(1 - \sqrt{1 - \left(\frac{\delta}{\mu}\right)^2}\right). \end{aligned} \tag{5.18}$$

A straightforward computation shows $\delta/\mu \leq \sqrt{2\sqrt{2} - 2}$ is equivalent to

$$1 - \sqrt{1 - \left(\frac{\delta}{\mu}\right)^2} \leq \frac{1}{\sqrt{2}} \left(\frac{\delta}{\mu}\right)^2. \tag{5.19}$$

Thus

$$d(x,y) \leq \mu \left(\frac{\delta}{\mu}\right)^2 \leq \tau\delta.$$

For the case $x = z$ it is apparent that (5.8) holds with $\delta = \mu$.

Finally, for the case $x \in \partial_{\mathrm{rel}} F_{\alpha'}(z,\mu)$ it is easy to see that (5.9) holds with $\delta = 0$. □

The following lemma is an immediate corollary of the first part of the preceding one where z is replaced by z', x is replaced by z, μ is replaced by ε', and ε is taken to be a number not necessarily equal to δ. In fact, Case I corresponds to $\varepsilon > \delta$, Case II to $\varepsilon = \delta$, and Case III to either $\varepsilon < \delta$ or an empty intersection.

LEMMA 5.4. *Let $S_{\alpha'} \subset S_\alpha$. Given $z \in S_\alpha$, and $z' \in S_{\alpha'}$, one of the following mutually exclusive relations hold:*

(I) $S_\alpha^\perp(z) \cap F_{\alpha'}(z',\varepsilon') \supsetneq F_\alpha(z,\varepsilon)$;
(II) $S_\alpha^\perp(z) \cap F_{\alpha'}(z',\varepsilon') = F_\alpha(z,\varepsilon)$;
(III) $S_\alpha^\perp(z) \cap F_{\alpha'}(z',\varepsilon') \subsetneq F_\alpha(z,\varepsilon)$.

Remark 5.1. Observe that by the Pythagorean theorem $d(z, \partial_{\mathrm{rel}} F_{\alpha'}(z,\mu))) > \mu$.

The authors found the schematic representations of Figures 6 and 7 useful mental aids for the arguments to follow.

Definition 5.1. (The set Ω) With respect to an n-tuple $\boldsymbol{\varepsilon} = (\varepsilon_0, \varepsilon_1, \ldots, \varepsilon_{n-1})$ of positive numbers set

$$\mathbf{T}_j = \mathbf{T}_j(\varepsilon_j) = \bigcup_{\dim S_\alpha = j} T_{\varepsilon_j}(S_\alpha),$$

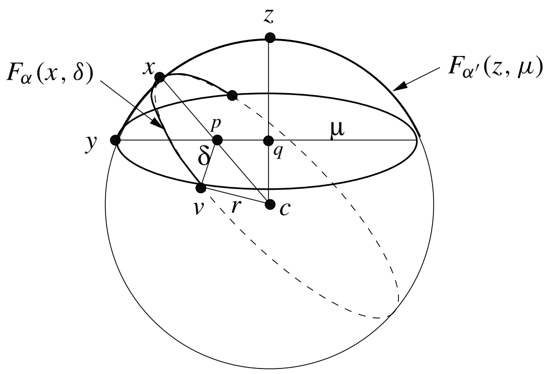

FIGURE 6. Spherical cap intersecting a great 2-sphere.

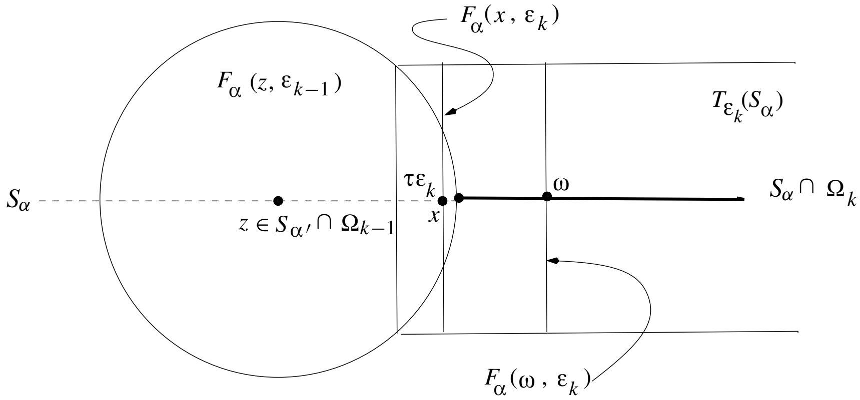

FIGURE 7. Schematic aid.

for $j = 0, 1, \ldots, n-1$. Next set

$$\Omega_0 = \Big(\bigcup_{\dim S_\alpha = 0} S_\alpha\Big)$$

and

$$\Omega_k = \Omega_k(\varepsilon_0, \ldots, \varepsilon_{k-1}) = \Big(\bigcup_{\dim S_\alpha = k} S_\alpha\Big) \Big\backslash \bigcup_{j<k} \mathbf{T}_j(\varepsilon_j),$$

for $k = 1, \ldots, n$. Note that sets in this hierarchy may be empty because there are not enough affine subspaces in $\mathcal{V}$ or the ε_j are too big. Finally, set

$$\Omega = \Omega(\boldsymbol{\varepsilon}) = \bigcup_{k \le n} \Omega_k.$$

Definition 5.2. An n-tuple $\boldsymbol{\varepsilon} = (\varepsilon_0, \varepsilon_1, \ldots, \varepsilon_{\dim S - 1})$ of positive numbers is said to *decrease sufficiently rapidly and prevent improper tube interference* with respect to a stratification $(S, \mathcal{V})$ if

(1) for $k > 0$, $\varepsilon_k/\varepsilon_{k-1} \le \tau$, with $\varepsilon_0 \le \tau r$, where $\tau = \sqrt{2 - 2\sqrt{2/3}}/8$;
(2) for $0 \le k \le \dim S - 1$ and $\omega \in \Omega_k$, if $V_\alpha \subset V_i \notin \mathcal{V}_\omega$, then $U_{8\varepsilon_k}(\omega) \cap T_{\varepsilon_k}(S_\alpha) = \emptyset$.

Item (1) will ensure that the ε_k's decrease fast enough. The value of τ is chosen so that

$$U_{8\varepsilon_k}(z) \subset U_{(\sqrt{2-2\sqrt{2/3}})\varepsilon_{k-1}}(z). \tag{5.20}$$

This relation implies that a $8\varepsilon_k$-neighborhood of any point in a sphere of radius ε_{k-1} is contained in an octant of that sphere centered around the given point. In §5.4.2, claim 2, of the proof of the Main Lemma we shall need that the hypotenuse of a geodesic spherical right triangle is longer than the other sides. This will be the case if the triangle lies in a neighborhood contained in an octant. In addition this value of τ is smaller than the quantity $\sqrt{2\sqrt{2} - 2}$ which appears in Lemma 5.3. It is also small enough that each time a geodesic belongs to the discussion (e.g. see §5.4.3, claim 3, in the proof of the Main Lemma) it is unique.

Item (2), which might impose that ratios of successive ε_k are much smaller than τ, guarantees that spherical caps fibering tubes in general do not interfere in unexpected ways with spherical caps centered on points of Ω and leaves enough room to simplify some arguments. We call attention of the reader to the fact that the subscript k in $T_{\varepsilon_k}(S_\alpha)$ is not necessarily the same dimension as that of S_α as it would be in the definition of $\mathbf{T}_k$ in Definition 5.1.

LEMMA 5.5. *With respect to the stratification* $(S, \mathcal{V})$ *let* $\boldsymbol{\varepsilon} = (\varepsilon_0, \varepsilon_1, \ldots, \varepsilon_{n-1})$ *decrease sufficiently rapidly so as to prevent improper tube interference. Then for each* $x \in \bigcup_{j<n} \mathbf{T}_j$ *there exist a largest integer* $0 \le k(x) \le n-1$*, a great sphere* S_α *with* $\dim S_\alpha = k(x)$*, and a point* $z \in S_\alpha$ *such that the* $(\alpha, \varepsilon_{k(x)})$*-spherical cap* $F_\alpha(z, \varepsilon_{k(x)})$ *contains* x *and satisfies*

$$\partial_{\mathrm{rel}} F_\alpha(z, \varepsilon_{k(x)}) \cap (\mathbf{T}_{-1} \cup \mathbf{T}_1 \cup \cdots \cup \mathbf{T}_{k(x)-1}) = \emptyset,$$

where $\mathbf{T}_{-1} = \emptyset$*.*

Moreover, this $(\alpha, \varepsilon_{k(x)})$*-spherical cap with top at* z *and containing* x *is unique and will be denoted by* $F_{\alpha(x)}(z(x), \varepsilon_{k(x)})$ *to show all the dependencies on* x*. However, for the sake of brevity we shall also denote it simply by* $F_\alpha(z, \varepsilon_{k(x)})$*. In the abbreviated form it is to be understood that the dependence of* α *and* z *on* x *is adequately indicated by the subscript* $k(x)$ *on* ε*. Here the unique great sphere* S_α *containing* z *determined by* x *we denote by* $S_{\alpha(x)}$*.*

Proof. Take the largest $k_1 < n$ such that $x \in \mathbf{T}_{k_1}$. By (5.4) x belongs to some spherical cap $F_{\alpha_1}(z_1, \varepsilon_{k_1})$ where $\dim S_{\alpha_1} = k_1$. If $\partial_{\mathrm{rel}} F_{\alpha_1}(z_1, \varepsilon_{k_1}) \cap \mathbf{T}_j = \emptyset$ for all $j < k_1$ we are done and the spherical cap $F_{\alpha_1}(z_1, \varepsilon_{k_1})$ is unique.

The argument for the uniqueness is the following. There are two cases:
(1) $z_1 \in \Omega_{k_1} \cap T_{\varepsilon_{k_1}}(S_{\alpha_1})$;
(2) $z_1 \in T_{\varepsilon_j}(S_{\alpha'})$, $j < k_1$, $\dim S_{\alpha'} = j$ and $S_{\alpha'} \subset S_{\alpha_1}$.
In case (1) $F_{\alpha_1}(z_1, \varepsilon_{k_1}) \subset U_{8\varepsilon_{k_1}}(z_1)$. The uniqueness follows from the fact that there are no improper tube interferences.

In case (2) we have $z_1 \in T_{\varepsilon_j}(S_{\alpha'})$ with $j < k_1$, say $z_1 \in F_{\alpha'}(z, \varepsilon_j)$. Lemma 5.3 implies that $d(z_1, \omega) \le \tau\varepsilon_{k_1}$ for some $\omega \in \Omega_{k_1}$ (in that lemma the role of x is played by z_1, δ by ε_{k_1}, μ by ε_j, and z by itself). Here $F_{\alpha_1}(z_1, \varepsilon_{k_1}) \subset U_{8\varepsilon_{k_1}}(\omega)$.

Hence also in case (2), the uniqueness follows by the way ε_{k_1} is chosen so that $F_{\alpha_1}(z_1, \varepsilon_{k_1})$ does not intersect any other $(\varepsilon_{k_1}, \alpha'')$-spherical cap with $\dim S_{\alpha''} = k_1$.

If $\partial_{\mathrm{rel}} F_{\alpha_1}(z_1, \varepsilon_{k_1}) \cap \mathbf{T}_j \neq \emptyset$ for some $k_2 < k_1$, take the largest such k_2. By Lemma 5.4 and the definition of ε_{k_1} we get that $x \in F_{\alpha_1}(z_1, \varepsilon_{k_1}) \subset \mathbf{T}_{k_2}$, so that x belongs to a spherical cap $F_{\alpha_2}(z_2, \varepsilon_{k_2})$ where $\dim S_{\alpha_2} = k_2$. Either we are done and this cap is unique or we must repeat the procedure finding k_3 and so forth. This procedure terminates at some $k \geq 0$ which defines $k(x)$ and a unique spherical cap $F_\alpha(z, \varepsilon_{k(x)})$ satisfying the conditions of this lemma. □

MAIN LEMMA 5.6. *Let $(S, \mathcal{V})$ be a stratification with* $\dim S = n$ *and* $\Omega(\boldsymbol{\varepsilon}) = \Omega(\varepsilon_0, \ldots, \varepsilon_{n-1})$ *as in Definition 5.1. If $\boldsymbol{\varepsilon} = (\varepsilon_0, \varepsilon_1, \ldots, \varepsilon_{n-1})$ decreases sufficiently rapidly and prevents improper tube interference, then*

$$\mathrm{clos}(x, \Omega) \subset \overline{\sigma(x)}.$$

Before giving a proof of this lemma, we need some more notation.

5.3. *Induced stratification and related constructs.* Choose some α and set $\dim V_\alpha = k + 1$, hence $\dim S_\alpha = k$. Pick any $x \in S_\alpha$, and set $\widehat{S} = \partial_{\mathrm{rel}} F_\alpha(x, \varepsilon_k)$ and $\widehat{\mathcal{V}} = \mathcal{V}_x = \{V_i : i \in \alpha\}$. Recall (see (5.3)) that the sphere $\widehat{S}$ has dimension $n - k - 1$. We denote the stratification induced by $\widehat{\mathcal{V}}$ on $\partial_{\mathrm{rel}} F_\alpha(x, \varepsilon_k)$ by $\widehat{\mathbb{S}} = (\widehat{S}, \widehat{\mathcal{V}})$. Let

$$\begin{aligned} \widehat{\boldsymbol{\varepsilon}} &= (\widehat{\varepsilon}_0, \ldots, \widehat{\varepsilon}_{n-k-2}) \\ &= (\varepsilon_{k+1}, \ldots, \varepsilon_{n+k-1}). \end{aligned} \tag{5.21}$$

For $\breve{\alpha} \subset \alpha$ we denote the great sphere $\widehat{S} \cap V_{\breve{\alpha}}$ by $\widehat{S}_{\breve{\alpha}}$ and the ε-tubular neighborhood $(V_{\breve{\alpha}} \times V_{\breve{\alpha}}^{\perp}(\varepsilon)) \cap \widehat{S}$ of $\widehat{S}_{\breve{\alpha}}$ in $\widehat{S}$ by $\widehat{T}_\varepsilon(\widehat{S}_{\breve{\alpha}})$. Notice that

$$\widehat{S}_{\breve{\alpha}} = S_{\breve{\alpha}} \cap \widehat{S}, \tag{5.22}$$

and that by (5.1)

$$\widehat{T}_\varepsilon(\widehat{S}_{\breve{\alpha}}) = T_\varepsilon(S_{\breve{\alpha}}) \cap \widehat{S}. \tag{5.23}$$

Furthermore, one can check that the cells $\widehat{\sigma}(x)$ of the induced stratification are obtained by intersecting the original cells with $\widehat{S}$: i.e.

$$\widehat{\sigma}(x) = \sigma(x) \cap \widehat{S}. \tag{5.24}$$

Next set

$$\widehat{\mathbf{T}}_j(\widehat{\varepsilon}_j) = \bigcup_{\substack{\dim \widehat{S}_{\breve{\alpha}} = j \\ \breve{\alpha} \subset \alpha}} \widehat{T}_{\widehat{\varepsilon}_j}(\widehat{S}_{\breve{\alpha}}), \tag{5.25}$$

for $j = 0, 1, \ldots, n - k - 2$, where by (5.21) and (5.23):

$$\widehat{T}_{\widehat{\varepsilon}_j}(\widehat{S}_{\breve{\alpha}}) = T_{\varepsilon_{j+k+1}}(S_{\breve{\alpha}}) \cap \widehat{S}.$$

From the obvious relation

$$\dim V_{\breve{\alpha}} = \dim V_\alpha + \dim(V_{\breve{\alpha}} \cap V_\alpha^{\perp})$$

and the definitions of $\widehat{S}$ and of j (as $\dim \widehat{S}_{\breve{\alpha}}$), we get $\dim V_{\breve{\alpha}} = j + k + 2$, thus

$$\begin{aligned}\widehat{\mathbf{T}}_j(\widehat{\varepsilon}_j) &= \bigcup_{\substack{\dim S_{\breve{\alpha}} = j+k+1 \\ \breve{\alpha} \subset \alpha}} T_{\varepsilon_{j+k+1}}(S_{\breve{\alpha}}) \cap \widehat{S} \\ &= \mathbf{T}_{j+k+1}(\varepsilon_{j+k+1}) \cap \widehat{S}. \end{aligned} \tag{5.26}$$

Next set

$$\widehat{\Omega}_0 = \bigcup_{\substack{\dim \widehat{S}_{\breve{\alpha}} = 0 \\ \breve{\alpha} \subset \alpha}} \widehat{S}_{\breve{\alpha}}$$

and

$$\widehat{\Omega}_m = \widehat{\Omega}_m(\widehat{\varepsilon}_0, \ldots, \widehat{\varepsilon}_{m-1}) = \Bigg(\bigcup_{\substack{\dim \widehat{S}_{\breve{\alpha}} = m \\ \breve{\alpha} \subset \alpha}} \widehat{S}_{\breve{\alpha}} \Bigg) \Bigg\backslash \bigcup_{j<m} \widehat{\mathbf{T}}_j(\widehat{\varepsilon}_j).$$

for $m = 1, \ldots, n - k - 1$. Finally, set

$$\widehat{\Omega} = \widehat{\Omega}(\boldsymbol{\varepsilon}) = \bigcup_{m \le n-k-1} \widehat{\Omega}_m.$$

From (5.22) and (5.26) we obtain the key relation:

$$\widehat{\Omega} = \Omega \cap \widehat{S}. \tag{5.27}$$

5.4. *Proof of Main Lemma 5.6.* We shall prove this by induction on the dimension of S.

INDUCTIVE HYPOTHESIS (n). *Let* $(S, \mathcal{V})$ *be a stratification with* $\dim S < n$. *If* $\boldsymbol{\varepsilon} = (\varepsilon_0, \varepsilon_1, \ldots, \varepsilon_{\dim S-1})$ *decreases sufficiently rapidly and prevents improper tube interference, then, for* $x \in S$,

$$\operatorname{clos}(x, \Omega) \subset \overline{\sigma(x)}.$$

After some preliminary discussion, the proof will consist in proving four claims which localize $\operatorname{clos}(x, \Omega)$ with greater and greater precision. The fourth and final claim finishes the induction and hence the proof of the Main Lemma 5.6.

Repeating here the contents of Definition 5.2 for the convenience of the reader, the premise of the inductive hypothesis states that:

(1) for $m > 0$, $\varepsilon_m/\varepsilon_{m-1} \le \tau$, with $\varepsilon_0 \le \tau r$, where $\tau = \sqrt{2 - 2\sqrt{2/3}}/8$;

(2) for $0 \le m \le \dim S - 1$ and $\omega \in \Omega_m$, if $V_\alpha \subset V_i \notin \mathcal{V}_\omega$, then $U_{8\varepsilon_m}(\omega) \cap T_{\varepsilon_m}(S_\alpha) = \emptyset$.

For a stratification $(\widehat{S}, \widehat{\mathcal{V}})$ induced by the stratification $(S, \mathcal{V})$ we get that $\widehat{\varepsilon}_m$ inherits property (1) from ε_{m+k+1}. Also $U_{8\widehat{\varepsilon}_m}(\widehat{\omega})$ inherits property (2) from $U_{8\varepsilon_{m+k+1}}(\widehat{\omega})$: i.e. $U_{8\widehat{\varepsilon}_m}(\widehat{\omega}) \cap \widehat{T}_{\widehat{\varepsilon}_m}(\widehat{S_{\widehat{\alpha}}}) = \emptyset$, since $U_{8\varepsilon_{m+k+1}}(\widehat{\omega}) \cap T_{\varepsilon_{m+k+1}}(S_{\widehat{\alpha}}) = \emptyset$ whenever $V_{\widehat{\alpha}} \subset V_i \notin \mathcal{V}_{\widehat{\omega}}$.

Under the assumption that the premise of the inductive hypothesis holds for the stratification $(S, \mathcal{V})$ where $\dim S < n$, we get from the above inheritance properties that the premise holds for all induced stratifications $(\widehat{S}, \widehat{\mathcal{V}})$ with $\dim \widehat{S} < n - 1$. So the conclusion of the inductive hypothesis then holds for these lower dimensions. This fact will come into play in the final claim.

To start the induction at $n = 1$, observe that when $\dim S = 0$, S consists of two points and the induction hypothesis (1) immediately follows.

Next, assume that the inductive hypothesis (n) holds. To prove the induction hypothesis ($n + 1$) let S be a sphere $\dim S = n$ and $\boldsymbol{\varepsilon} = (\varepsilon_0, \varepsilon_1, \ldots, \varepsilon_{\dim S-1})$ decrease sufficiently rapidly and prevent improper tube interference.

Take $x \in S$. If $x \in \Omega$ there is nothing to prove. Suppose $x \notin \Omega$. Then $x \in \bigcup_{j<n} \mathbf{T}_j$, and we get a unique spherical cap $F_\alpha(z, \varepsilon_{k(x)})$ which satisfies Lemma 5.5. Recall equation (5.2), i.e.

$$P_\alpha(z, 3\varepsilon_{k(x)}) = \bigcup_{y \in U_{3\varepsilon_{k(x)}}(z) \cap \sigma(z)} F_\alpha(y, 3\varepsilon_{k(x)}).$$

(In this expression, as in the convention for spherical caps, the dependence of α and z on x is indicated by the subscript $k(x)$ on ε.) Either $z = \omega \in \Omega_{k(x)}$, or $z \notin \Omega_{k(x)}$ and there exists $\omega \in \Omega_{k(x)}$ such that $d(z, \omega) < \tau\varepsilon_{k(x)}$ by using Lemma 5.3. In either case, by (5.6), we have

$$P_\alpha(z, 3\varepsilon_{k(x)}) \subset U_{8\varepsilon_{k(x)}}(\omega). \tag{5.28}$$

5.4.1.

CLAIM 1. $\operatorname{clos}(x, \Omega) \subset P_\alpha(z, 3\varepsilon_{k(x)})$.

Proof. Due to the definition of $k(x)$ and property (2) of Definition 5.2, the sphere $\partial_{\mathrm{rel}} F_\alpha(z, \varepsilon_{k(x)})$ has a non-empty intersection with Ω: so the distance from $x \in F_\alpha(z, \varepsilon_{k(x)})$ to Ω is less than or equal to the diameter$(F_\alpha(z, \varepsilon_{k(x)})) = 2\varepsilon_{k(x)}$. Claim 1 follows because $d(x, S \setminus P_\alpha(z, 3\varepsilon_{k(x)})) > 2\varepsilon_{k(x)}$. □

5.4.2.

CLAIM 2. $\operatorname{clos}(x, \Omega) \subset F_\alpha(z, 3\varepsilon_{k(x)}) \cup (U_{3\varepsilon_{k(x)}}(z) \cap \sigma(z))$.

Proof. Let $\omega \in \operatorname{clos}(x, \Omega)$. We know by Claim 1 that $\omega \in P_\alpha(z, 3\varepsilon_{k(x)})$: so we can assume that $\omega \in P_\alpha(z, 3\varepsilon_{k(x)}) \setminus U_{3\varepsilon_{k(x)}}(z) \cap \sigma(z)$. We just have to prove that $\omega \in F_\alpha(z, 3\varepsilon_{k(x)})$. From Claim 1 there exists $y \in U_{3\varepsilon_{k(x)}}(z) \cap \sigma(z)$ such that $\omega \in F_\alpha(y, 3\varepsilon_{k(x)}) \setminus U_{3\varepsilon_{k(x)}}(z) \cap \sigma(z)$. Project the point ω to $F_\alpha(z, 3\varepsilon_{k(x)})$ by using the isometry given in Lemma 5.2, say

$$\omega' = O^{-1}(\omega) \in F_\alpha(z, 3\varepsilon_{k(x)}) \setminus F_\alpha(z, \varepsilon_{k(x)}).$$

Now there are two possibilities.

$$\text{Either} \quad \omega' \in \Omega \quad \text{or} \quad \omega' \in \mathbf{T}_j, \quad j < k(x).$$

This last possibility does not occur because it would imply

$$\partial_{\mathrm{rel}} F_\alpha(z, \varepsilon_{k(x)}) \subset \mathbf{T}_j,$$

by using Lemma 5.3, contradicting the definition of $k(x)$ in Lemma 5.5. Hence, $\omega' \in \Omega$.

Assume $\omega' \neq \omega$. Then, by the property of geodesic right triangles resulting from Definition 5.2(1) we get $|\omega - x| > |\omega' - x|$ which contradicts the fact that $\omega \in \operatorname{clos}(x, \Omega)$. Hence $\omega = \omega'$. □

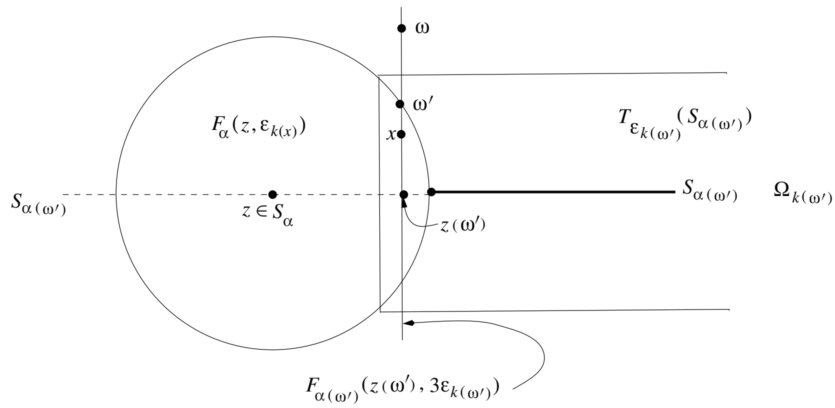

FIGURE 8. Schematic aid to Claim 3. Here the disk corresponds to a transversal view of the horizontal strip in Figure 7.

5.4.3.

CLAIM 3. $\mathrm{clos}(x, \Omega) \subset \partial_{\mathrm{rel}} F_\alpha(z, \varepsilon_{k(x)}) \cup (U_{3\varepsilon_{k(x)}}(z) \cap \sigma(z))$.

Proof. Let $\omega \in \mathrm{clos}(x, \Omega)$. We know by Claim 2 that $\omega \in F_\alpha(z, 3\varepsilon_{k(x)}) \cup (U_{3\varepsilon_{k(x)}}(z) \cap \sigma(z))$, so we can assume $\omega \in F_\alpha(z, 3\varepsilon_{k(x)}) \setminus \{z\}$. Along the great circle on S from ω through x we have a point $\omega' \in \partial_{\mathrm{rel}} F_\alpha(z, \varepsilon_{k(x)})$ which is well defined by taking the shortest geodesic from ω to ω'. Again there are two possibilities.

$$\text{Either} \quad \omega' \in \Omega \quad \text{or} \quad \omega' \in \mathbf{T}_j, \quad \text{for some } j > k(x).$$

The second possibility does not occur for the following reason. Using Lemma 5.5 we can take $j = k(\omega') > k(x)$. This means that $\omega' \in F_{\alpha(\omega')}(z(\omega'), \varepsilon_{k(\omega')})$ for some $z(\omega') \in S_{\alpha(\omega')} \supset S_\alpha$ with $\dim S_{\alpha(\omega')} = k(\omega')$ and $\partial_{\mathrm{rel}} F(z(\omega'), \varepsilon_{k(\omega')}) \cap (\mathbf{T}_0 \cup \cdots \cup \mathbf{T}_{k(\omega')-1}) = \emptyset$. Also, we would have that $\omega \in \mathrm{clos}(\omega', \Omega)$, for otherwise $\omega \notin \mathrm{clos}(x, \Omega)$. Using Claim 2, with ω' playing the role of x in that claim, we would get $\omega \in F_{\alpha(\omega')}(z(\omega'), 3\varepsilon_{k(\omega')})$. Since $\omega, \omega' \in S^\perp_{\alpha(\omega')}(z(\omega'))$, we would also get $x \in S^\perp_{\alpha(\omega')}(z(\omega'))$. Applying Lemma 5.4 Case III, we would get $x \in F_{\alpha(\omega')}(z(\omega'), \varepsilon_{k(\omega')})$ (see Figure 8) which would imply that $k(x) \geq k(\omega')$, a contradiction completing the proof of Claim 3. □

5.4.4.

CLAIM 4. $\mathrm{clos}(x, \Omega) \subset \overline{\sigma(x)}$.

Proof. From the previous claim

$$\mathrm{clos}(x, \Omega) \subset \partial_{\mathrm{rel}} F_\alpha(z, \varepsilon_{k(x)}) \cup U_{3\varepsilon_{k(x)}}(z) \cap \sigma(z).$$

If $x = z$, then, by virtue of (5.10) and the fact that the distance from x to the boundary of the $\varepsilon_{k(x)}$-tubular neighborhood of z is greater than $\varepsilon_{k(x)}$ (see Remark 5.1) we get,

$$\mathrm{clos}(x, \Omega) \subset U_{3\varepsilon_{k(x)}}(z) \cap \overline{\sigma(x)}.$$

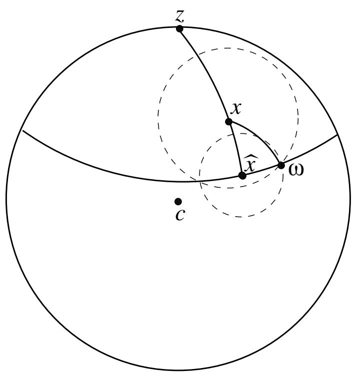

FIGURE 9. Reduction to a great 2-sphere in three dimensions.

So assume $x \neq z$ and

$$\operatorname{clos}(x, \Omega) \cap \partial_{\mathrm{rel}} F_\alpha(z, \varepsilon_{k(x)}) \setminus U_{3\varepsilon_{k(x)}}(z) \cap \sigma(z) \neq \emptyset. \tag{5.29}$$

Project $x \in F_\alpha(z, \varepsilon_{k(x)})$ along the great circle on S from z through x to $\widehat{x} \in \partial_{\mathrm{rel}} F_\alpha(z, \varepsilon_{k(x)})$. Since x and $\widehat{x}$ lie on an arc of a great circle which does not cross cell boundaries, $\widehat{x} \in \overline{\sigma(x)}$. Hence

$$\overline{\sigma(x)} = \overline{\sigma(\widehat{x})}.$$

If $\widehat{x} \in \Omega$, then $\operatorname{clos}(x, \Omega) = \{\hat{x}\}$, and we are finished. If $\widehat{x} \notin \Omega$, we use the inductive hypothesis to get

$$\operatorname{clos}(\widehat{x}, \widehat{\Omega}) \subset \overline{\widehat{\sigma}(\widehat{x})}.$$

We can assert for $\widehat{x} \in \partial_{\mathrm{rel}} F_\alpha(z, \varepsilon_{k(x)})$ that

$$\overline{\widehat{\sigma}(\widehat{x})} = \overline{\sigma(\widehat{x})} \cap \partial_{\mathrm{rel}} F_\alpha(z, \varepsilon_{k(x)}) \subset \overline{\sigma(\widehat{x})} \tag{5.30}$$

despite the fact that it does not hold for spherical cap boundaries in general. The reason for this is that affine subspaces in $\mathcal{V}$ which intersect strata of $\widehat{S} = \partial_{\mathrm{rel}} F_\alpha(z, \varepsilon_{k(x)})$ to form cells in $\widehat{S}$ are precisely those in $\mathcal{V}$ which contain V_α. From (5.27) and (5.29)

$$\operatorname{clos}(x, \Omega) = \operatorname{clos}(x, \widehat{\Omega}).$$

However, we have

$$\operatorname{clos}(x, \widehat{\Omega}) = \operatorname{clos}(\widehat{x}, \widehat{\Omega}),$$

the reason for which is the following.

Observe that

$$\widehat{S} \cap S(x, d(x, \Omega)) = \widehat{S} \cap S(\widehat{x}, d(\widehat{x}, \widehat{\Omega}))$$

as illustrated in Figure 9 where we have drawn the great 2-sphere in three dimensions determined by the center c and the points $z, x, \omega \in \operatorname{clos}(x, \widehat{\Omega})$. $\widehat{S}$ intersects this sphere in the meridian containing $\widehat{x}$ and ω. Thus,

$$\begin{aligned}\operatorname{clos}(x, \widehat{\Omega}) &= (\widehat{S} \cap S(x, d(x, \Omega))) \cap \widehat{\Omega} \\ &= (\widehat{S} \cap S(\widehat{x}, d(\widehat{x}, \widehat{\Omega}))) \cap \widehat{\Omega} \\ &= \operatorname{clos}(\widehat{x}, \widehat{\Omega}).\end{aligned}$$

This finishes the proof of Claim 4 and also the induction step in the proof of the Main Lemma 5.6. □

□

5.5. *Conclusion of proof of Proposition 4.1.* Let $S = S(1, 0)$ be the $(N-1)$-dimensional unit sphere centered at the origin and $\mathcal{V}$ be the set of codimension-1 subspaces of $\mathbb{R}^N$ normal to the edges of the polytope P.

5.5.1. *Proof of Proposition 4.1(3).* The closest points $\omega \in \Omega$ to $y \in \bar{\sigma} \subset C_1(v_i)$ are in the same cell closure $\bar{\sigma}$. Each $C(v_i)$ is a polyhedral cone. The half-spaces, of which it is the intersection, are determined by codimension-one subspaces orthogonal to the edges connecting v_i to neighboring vertices. Let $\mathcal{V}_P$ be the collection of codimension-one subspaces orthogonal to all the edges emanating from all the vertices of P. Let the collection $\mathcal{V}$ of codimension-one subspaces of §5 contain $\mathcal{V}_P$. The Main Lemma 5.6 holds for either collection of codimension-one subspaces because the partition of S into cells σ resulting from $\mathcal{V}$ merely refines the one resulting from $\mathcal{V}_P$. That is to say that the cell containing $\operatorname{clos}(x, \Omega)$ given by $\mathcal{V}$ is a subset of the cell containing $\operatorname{clos}(x, \Omega)$ given by $\mathcal{V}_P$. This fact is the basis of the proof of Theorem 6.2.

All that is left to show is that there exists $\boldsymbol{\varepsilon} = (\varepsilon_0, \varepsilon_1, \dots, \varepsilon_{n-1})$ which satisfies the premise of the Main Lemma 5.6: namely, $\boldsymbol{\varepsilon}$ decreases sufficiently rapidly and prevents improper tube interference. Upon establishing this, the conclusion of the Main Lemma 5.6 will then hold.

Let $x \in S$. Because $\mathcal{V}$ is a finite collection, the planes in $\mathcal{V}$ which do not pass through x will have a certain minimal distance to x. Hence, a small enough neighborhood of x avoids all planes that do not contain x; and from the definition of $\sigma(x)$, we then get that for every $x \in S$ there is $\varepsilon(x) > 0$ such that the following holds:

(1) for all $y \in U_{\varepsilon(x)}(x) \cap \sigma(x)$, $\mathcal{V}_y = \mathcal{V}_x$;

(2) if $V_\alpha \subset V_i \notin \mathcal{V}_x$, then $U_{8\varepsilon(x)}(x) \cap T_{\varepsilon(x)}(S_\alpha) = \emptyset$.

By virtue of compactness of the various components of $\Omega_k(\varepsilon_0, \dots, \varepsilon_{k-1})$, we can then inductively choose a uniform ε_k for Ω_k so that:

(1) for $k > 0$, $\varepsilon_k/\varepsilon_{k-1} \leq \tau$, with $\varepsilon_0 \leq \tau r$, where $\tau = \sqrt{2 - 2\sqrt{2/3}}/8$;

(2) for every $\omega \in \Omega_k$ if $V_\alpha \subset V_i \notin \mathcal{V}_\omega$, then $U_{8\varepsilon_k}(\omega) \cap T_{\varepsilon_k}(S_\alpha) = \emptyset$,

as needed. □

5.5.2. *Proof of Proposition 4.1(2).* For each cell $\sigma \subset C_1(v_i)$ there is a $\varepsilon_\sigma > 0$ such that $\Omega \cap U_{\varepsilon_\sigma}(\bar{\sigma}) \subset \bar{\sigma}$. Since $d(z, \partial_{\mathrm{rel}} F_\alpha(z, \varepsilon_k)) > \varepsilon_k$, we can choose $\varepsilon_\sigma = \varepsilon_k$, where k is the dimension of σ. □

5.5.3. *Proof of Proposition 4.1(1).* Ω is not confined to any half-sphere. There are several ways of proving this. One way is to choose ε_k so that the total volume of all ε-tubular neighborhoods is less than half the measure of S. Another is to use the fact that $\mathcal{V}$ is the set of codimension-one subspaces normal to edges of P. Perhaps the easiest way is to choose n linearly independent vectors $v_1, \dots, v_N \in \Omega$ and observe that $-v_1, \dots, -v_N$ are also in Ω. □

□

Having completed the proof of Proposition 4.1 we have finished assembling all the elements that prove Theorem 1.1.

6. *Additional properties*

THEOREM 6.1. *The constructed Q is globally absorbing for the family $\{\phi_\gamma : \gamma \in P_0\}$ for any P_0 a compact subset of the interior of P, i.e. given a compact subset X of $\mathbb{R}^N$ there exists an n_0 (depending on P_0) such that if $n > n_0$ then $\phi_{\gamma_1} \cdots \phi_{\gamma_n}(X) \subset Q$ for $\gamma_1, \ldots, \gamma_n \in P_0$.*

We first remark that for a fixed compact subset P_0 of the interior of P any large enough sphere S centered at the origin will be a globally absorbing region: because any larger sphere contracts by a fixed amount under the action of any ϕ_γ, the amount of contraction depend on distance from P_0 to ∂P. However, the size of the sphere S varies with P_0. Theorem 6.1 is a stronger result. It states that Q can be so chosen that it is universally globally absorbing for all such P_0.

Proof. Let $\rho_0 > R$ where R is given by (3.10). We shall show that the set $Q = \rho_0 Q_\infty$ is globally absorbing. From the strengthening of (3.5) to strict inequality which follows from a strengthening of (3.6) to strict inequality, the other participating inequality (3.9) already being strict, we get that for any $\gamma \in P_0$ and $\rho \geq \rho_0$ the set ρQ_∞ is mapped by ϕ_γ into $\rho' Q_\infty$ where $\rho - \rho' = \delta(\rho) > 0$ is a quantity which depends on the distance from P_0 to ∂P. The function $\delta(\rho)$ is increasing for $\rho \geq \rho_0$. Thus $\rho' < \rho - \delta(\rho_0)$, which leads to

$$\phi_{\gamma_1} \cdots \phi_{\gamma_n}(\rho Q) \subset (\rho - n\delta(\rho_0))Q \subset \rho_0 Q$$

for some n. □

THEOREM 6.2. *Given a finite set of polytopes and their time-dependent dynamical systems as defined in Theorem 1.1, there exists a Q_∞ such that the set $Q = \rho Q_\infty$ is invariant for large enough ρ for each dynamical system.*

Proof. By virtue of the remark in §5.5.1 we just take $\mathcal{V}$ to contain all the codimension-one subspaces orthogonal to all the one-dimensional edges of all the polytopes. □

Acknowledgements. Shub's work was partly funded by NSF Grant #DMS-9988809. Martens' and Tresser's work was partly funded by NSF Grant#DMS-0073069.

REFERENCES

[1] R. L. Adler, B. Kitchens, M. Martens, A. Nogueira, C. Tresser and C. W. Wu. Error bounds for error diffusion and other mathematical problems arising in digital halftoning. *Proc. SPIE* **3963** (2000), 437–443.

[2] R. L. Adler, B. Kitchens, M. Martens, A. Nogueira, C. Tresser and C. W. Wu. *Error Bounds for Error Diffusion and Related Halftoning Algorithms (Proc. IEEE Int. Symp. on Circuits and Systems, 2001)*, pp. II-513–II-516.

[3] R. L. Adler, B. P. Kitchens, M. Martens, C. P. Tresser and C. W. Wu. The mathematics of halftoning. *IBM J. Res. & Dev.* **47**(1) (2003), 1–9.

[4] R. Adler, B. Kitchens and C. Tresser. Dynamics of nonergodic piecewise affine maps of the torus. *Ergod. Th. & Dynam. Sys.* **21** (2001), 959–999.

[5] V. Baladi, D. Rockmore, N. Tongring and C. Tresser. Renormalization on the n-torus. *Nonlinearity* **5** (1992), 1111–1136.

[6] J. Beck and W. W. L. Chen. *Irregularities of Distribution*. Cambridge University Press, Cambridge, 1987.

[7] J. Bernoulli. Sur une nouvelle espèce de calcul. *Recueil Astron. (Berlin)* **1** (1772), 255–284.

[8] M. Broja, R. Eschbach and O. Bryngdahl. Stability of active binariation processes. *Opt. Comm.* **60** (1986), 353–358.

[9] E. B. Christoffel. Observatio arithmetica. *Ann. of Math.* (2) **6** (1875) 148–152.

[10] R. Floyd and L. Steinberg. An adaptive algorithm for spatial grey scale. *SID (Society for Information Display) Symposium, Digests of Papers*, (1975), pp. 36–37.

[11] J. F. Jarvis, C. N. Judice and W. H. Ninke. A survey of techniques for the display of continuous tone pictures on bilevel displays. *Computer Graphics Image Proc.* **5** (1976), 13–40.

[12] D. B. Judd and G. Wyszecki. *Color in Business, Science, and Industry*. Wiley, New York, 1975.

[13] H. Kang. *Color Technology for Electronic Imaging Devices*. SPIE Optical Engineering Press; Bellingham, WA, 1997.

[14] B. W. Koplatzik and C. A. Bouman. Optimized error diffusion for high quality image display. *J. Electronic Imaging* **1** (1992), 277–292.

[15] J. C. Lagarias and C. Tresser. A walk along the branches of the extended Farey tree. *IBM J. Res. & Dev.* **39** (1995), 283–294.

[16] M. Martens, T. Nowicki, M. Shub and C. Tresser. Convex dynamics: construction of invariant regions. In preparation.

[17] M. Martens, T. Nowicki, C. Tresser and C. W. Wu. Convex dynamics: bounds for scheduling. In preparation.

[18] A. A. Markoff. Sur une question de Jean Bernoulli. *Math. Ann.* **19** (1882), 27–36.

[19] H. G. Meijer. On a distribution problem in finite sets. *Nederl. Akad. Wetensch. Indag.* **35** (1973), 9–17.

[20] H. G. Meijer and H. Niederreiter. On a distribution problem in finite sets. *Compositio Math.* **25** (1972), 153–160.

[21] M. Morse and G. A. Hedlund. Symbolic dynamics II. Sturmian trajectories. *Amer. J. Math.* **62** (1940), 1–42.

[22] H. Niederreiter. On the existence of uniformly distributed sequences in compact spaces. *Compositio Math.* **25** (1972), 93–99.

[23] H. Niederreiter. A distribution problem in finite sets. *Applications of Number Theory to Numerical Analysis (Proc. Symp. University of Montréal, 1971)*. Ed. S. K. Zaremba. Academic Press, New York, 1972, pp. 237–248.

[24] T. Nowicki and C. Tresser. Convex dynamics: properties of invariant sets. *Nonlinearity* **17** (2004), 1645–1676.

[25] T. Nowicki and C. Tresser. Convex dynamics: unavoidable difficulties in bounding some greedy algorithms. *Chaos* **14** (2004), 55–71.

[26] R. T. Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, NJ, 1972.

[27] A. Schrijver. *Theory of Linear and Integer Programming (Wiley-Interscince Series in Discrete Mathematics and Optimization)*. Wiley, New York, 1996.

[28] R. M. Siegel, C. Tresser and G. Zettler. A decoding problem in dynamics and in number theory. *Chaos* **2** (1992), 473–493.

[29] H. J. S. Smith. Note on continued fractions. *Messenger Math.* **VI** (1877), 1–14.

[30] W. S. Stiles and G. Wyszecki. *Color Science: Concepts and Methods, Quantitative Data and Formulae*, 2nd edn. Wiley, New York, 1982.

[31] P. Stucki. MECCAA: Multiple-error correcting computation algorithm for bilevel image hard-copy reproduction. *Research Report RZ-1060*, IBM Research Laboratory, Zurich, Switzerland, 1981.

[32] J. R. Sullivan, R. L. Miller and T. J. Wetzel. Color digital halftoning with vector error diffusion. *US Patent* 5,070,413, 1991.

[33] R. Tijdeman. On a distribution problem in finite and countable sets. *J. Combin. Theory Ser. A* **15** (1973), 129–137.

[34] R. Tijdeman. The chairman assignment problem. *Discrete Math.* **32** (1980), 323–330.

[35] R. Ulichney. *Digital Halftoning*. MIT Press, Cambridge, MA, 1987.