Topology of Plane Sections of Periodic Polyhedra with an Application to the Truncated Octahedron

Roberto De Leo

CONTENTS

The main results of A. Zorich and I. Dynnikov regarding plane sections of periodic surfaces are extended to the piecewise linear case. As an application, the stereographic map of a truncated octahedron, extended to all of $\mathbb{R}^3$ by periodicity, is analyzed numerically.

1. INTRODUCTION

The problem of the asymptotics of plane sections of smooth periodic surfaces, extracted from the physics literature by S. P. Novikov in 1982 [Novikov 82], turned out to be much richer then expected, leading ultimately to the association of a fractal on $\mathbb{R}\,\mathrm{P}^2$ to every element of a large class of triply periodic functions in $\mathbb{R}^3$ (see Section 2).

In [De Leo 03] we developed a C++ library and used it to investigate the stereographic maps associated with triply periodic smooth functions in order to obtain information about the Novikov conjecture claiming that the Hausdorff dimension of the associated fractals is strictly between 1 and 2. Unfortunately, the running time of our numerical explorations grows unrealistically long as soon as we sample with resolutions fine enough to provide a hint of the fractals, mainly because the number of polygons of the meshes approximating a curved smooth surface (e.g., so as to determine its intersection with a plane) rapidly becomes very large.

This fact suggests that from the numerical point of view, polyhedra are the best surfaces to study, at the very least for the obvious reason that the number of polygons needed to describe them at any resolution is constant. Moreover, such constants can be rather small even in such nontrivial cases presented in this paper, as in the case of the *extended truncated octahedron*, which has just eight hexagonal faces.

2000 AMS Subject Classification: 57M50, 53D17, 53C12, 37E35, 65D18

Keywords: low-dimensional topology, Poisson geometry, foliations

1058-6458/2006 $0.50 per page

Since the main results of the theory, due to A. V. Zorich [Zorich 84] and I. A. Dynnikov [Dynnikov 97, Dynnikov 99], refer to the case of smooth surfaces only, in Section 2 we will provide independent proofs for those theorems that make use of the smooth structure and will mention the main properties of the system. Then, in Section 3, we present the algorithm we implemented to explore the problem numerically and our numerical results in the case of the polyhedron obtained by extending the truncated octahedron by periodicity. Note that the figures relative to highest-resolution numerical analysis can be found in the online appendix [De Leo 06] to this paper.

2. FUNDAMENTAL OBJECTS AND THEOREMS

2.1 Critical Points of Height Functions in Polyhedra

An analogue of the Morse theory for height functions on polyhedra was introduced by T. Banchoff in [Banchoff 67, Banchoff 70]. Here we recall the concepts relevant for the present paper, slightly modified to cover the case of periodic polyhedra.

Definition 2.1. By *embedded polyhedron* $M \subset \mathbb{R}^3$ we mean a countable collection of cells $K = \{C^r \subset \mathbb{R}^3\}_{r=0,1,2}$, where 0-cells are points (vertices), 1-cells are closed connected segments, and 2-cells are convex closed plane polygons such that the following hold:

P1. The boundary of any cell is the union of cells of lower dimension.

P2. Every cell having points in common with a cell of higher dimension is completely contained within it.

P3. K is locally finite, i.e., every vertex has a neighborhood in $\mathbb{R}^3$ that intersects only finitely many cells.

P4. For any point $p \in M$, the union $\mathrm{Star}(p)$ of all cells containing that point is homeomorphic to an open disk.

A triply periodic polyhedron is a polyhedron that is invariant with respect to a rank-3 discrete subgroup $\Gamma \simeq \mathbb{Z}^3$ of $\mathbb{R}^3$. Finally, by the polyhedron $\mathcal{M} \subset \mathbb{T}^3$ we mean the quotient $M/\Gamma \subset \mathbb{R}^3/\Gamma \simeq \mathbb{T}^3$ of a triply periodic polyhedron.

Every triply periodic polyhedron M embedded in $\mathbb{R}^3$ is the lift of a compact polyhedron $\mathcal{M}$ embedded in $\mathbb{T}^3$. Since we are going to study plane foliations of polyhedra, we are interested in the reciprocal relation between polyhedra and height functions (or equivalently, constant 1-forms):

Definition 2.2. A height function $h(p) = h^\alpha p_\alpha$ is called *generic* for the polyhedron M if no edge of M is perpendicular to the direction $\boldsymbol{H} = (h^\alpha)$. Equivalently, a constant 1-form $\omega = h^\alpha dp_\alpha$ in $\mathbb{R}^3$ (respectively $\mathbb{T}^3$) is generic for M (respectively $\mathcal{M}$) if no edge of M (respectively $\mathcal{M}$) is contained in a single leaf of ω.[1]

Since height functions are not single-valued on $\mathbb{T}^3$ (unless $\boldsymbol{H}$ is an integer direction, i.e., parallel to a lattice vector), while their differentials $\omega = dh$ are always well defined in both $\mathbb{R}^3$ and $\mathbb{T}^3$, we will refer mostly to 1-forms from now on. From the definition above it is clear that as in the smooth case, the set of nongeneric 1-forms has zero measure.

Note that the foliation induced on $\mathbb{T}^3$ (and therefore on $\mathcal{M}$) by ω does not change by multiplying the 1-form by a nonzero scalar, so from now on we will simultaneously think of ω as a constant 1-form and as a point in $\mathbb{R}\,\mathrm{P}^2$.

Definition 2.3. The index of a point $p \in M$ (or equivalently, $[p] \in \mathcal{M}$) with respect to a generic constant 1-form ω is the integer $i(p, \omega) = 1 - s/2$, where s is the number of segments having p as an endpoint in which the leaf of ω passing through p cuts $\mathrm{Star}(p)$. If $i(p, \omega) = 0$, i.e., if the $\mathrm{Star}(p)$ is cut in exactly two components, the point is said to be *regular*; otherwise, it is called *critical*.

The set of critical points for any generic constant 1-form is of course a subset of the set of vertices; exactly as in the smooth case, minima and maxima have index $+1$ and nondegenerate saddles have index -1. The main difference between the smooth and the piecewise linear (PL) case, for our purposes, is that saddles that are unstable in the former case, i.e., that disappear for small perturbations of the 1-form direction, are stable in the latter and therefore cannot be disregarded; the simplest example is provided by the "monkey saddle," which has index -2 (see Figure 1).

Nevertheless, an analogue of the critical point theorem for generic 1-forms holds:

[1] The Frobenius theorem shows that the distribution $\omega = 0$ is integrable iff the 1-form ω is closed. In this case, the leaves induced by $\omega = h^\alpha dp_\alpha$ in $\mathbb{T}^3$ are the projections of the $\mathbb{R}^3$ planes perpendicular to $\boldsymbol{H} = (h^\alpha)$; those induced on a polyhedron $\mathcal{M}$ are the intersections of these leaves with the polyhedron.

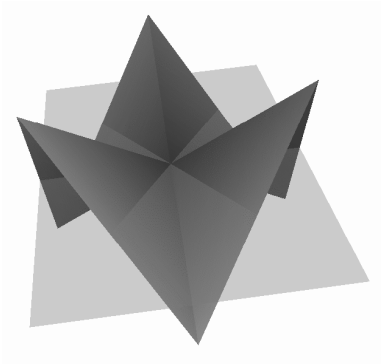

FIGURE 1. Piecewise linear monkey saddle: In the smooth case, an arbitrarily small perturbation would suffice to resolve this critical point into a pair of elementary saddles; in the piecewise linear case it is stable and therefore generic.

Theorem 2.4. *If M is a triply periodic polyhedron, invariant by the action of the rank-3 group $\Gamma \subset \mathbb{Z}^3$, and ω is generic for M, then $\sum_{[p]\in\mathcal{M}} i(p,\omega) = \chi(\mathcal{M})$, where $[p]$ is the set of vertices Γ-equivalent to p and the sum is hence extended to any set of inequivalent vertices. Equivalently, if we set the volume of a Dirichlet domain of Γ to 1, then the average of the Euler characteristic of M converges to the Euler characteristic of $\mathcal{M}$:*

$$\overline{\chi(M)} = \lim_{R\to\infty} \sum_{\|p\|<R} i(p,\omega)/\operatorname{Vol}(B_R) = \chi(\mathcal{M}).$$

Proof: Since Banchoff's proof [Banchoff 70] of the critical point theorem for polyhedra is based only on local identities that are trivially true in $\mathbb{T}^3$, that proof holds with no change in the case at hand.

The second part can be proved by considering that the genus g of the surface $\mathcal{M}$ contained inside a cube of side R can be evaluated by reducing the surface homotopically to a graph and then evaluating the rank of the graph's first homology group. The result follows from the consideration that every component contained in an inner unitary cube contributes g to the total genus of the component contained in the cube of radius R, and their number grows with R^3, while the cubes on the boundary provide a smaller contribution that can be disregarded in the limit for $R \to \infty$, since their number grows only as R^2. □

2.2 Structure of Foliations

In the most general case, a constant 1-form ω induces on a triply periodic polyhedron M both open and closed leaves.

Since being homotopic to zero is an open condition, leaves close enough to closed ones are also closed; maximal components of closed leaves are always enclosed between a pair of critical points of ω on M and form either cylinders (when both critical points are saddles) or disks (when one is a saddle and the other is a center) or spheres (when they are both centers). The last two cases are topologically trivial: A disk covered by closed leaves around a center is exactly a homotopy to a point of that component of M, and if M is a sphere, then no open orbit can ever be induced on it by a closed 1-form. Hence, once the topologically trivial components are removed, what is left is a collection of cylinders $\mathcal{C}_i$ that separates a collection of subpolyhedra with boundary $\mathcal{N}_i$ filled by open leaves.

Definition 2.5. A *genus-k component* of $\mathcal{M}$ is a piecewise linear submanifold with boundary $\mathcal{N}$ of $\mathcal{M}$ such that the following conditions are satisfied:

1. $\partial\mathcal{N}$ is the finite disjoint union of parallel plane (topological) circles homotopic to zero.
2. The closed polyhedron $\overline{\mathcal{N}}$ obtained by filling the holes of $\mathcal{N}$ with plane disks has genus k.

Definition 2.6. A *Dynnikov decomposition* $\mathcal{Z}$ of a polyhedron $\mathcal{M}$ is a collection of subpolyhedra with boundary $\{\mathcal{C}_i, \mathcal{N}_j\}$ of $\mathcal{M}$ such that the following conditions are satisfied:

1. Every $\mathcal{C}_i$ is homotopic either to a closed cylinder or to a closed disk.
2. Every $\mathcal{N}_j$ is a genus-k_j component of $\mathcal{M}$.
3. $i \neq j \implies \mathcal{N}_i \cap \mathcal{N}_j = \varnothing$, while every other pair of distinct subpolyhedra of $\mathcal{Z}$ that is not disjoint shares a single boundary component.
4. $\mathcal{M} = \bigcup_i \mathcal{C}_i \bigcup_j \mathcal{N}_j$.

The genus and rank of a Dynnikov decomposition $\mathcal{Z}$ are the highest genus and rank of the $\overline{\mathcal{N}}_j$'s contained in $\mathcal{Z}$. A genus-1 rank-2 Dynnikov decomposition is called a *Zorich decomposition*.

Remark 2.7. From the considerations above it is clear that every constant 1-form ω in general position with respect to $\mathcal{M}$, i.e., such that there are no saddle connections, naturally induces a Dynnikov decomposition $\mathcal{Z}$ on $\mathcal{M}$ in which all $\mathcal{C}_i$ are filled by the closed leaves and all $\mathcal{N}_j$ by the open ones.

Under the same assumptions, with every focus is associated a saddle that "cancels" it. Namely there is a homotopy of the surface that gets rid of the saddle-focus pair without modifying the topology of the open orbits nearby; we call the saddle-type critical points that are left *topological saddles*, since it is they that contribute to the Euler characteristic of the surface.

In the smooth generic case, the number of Γ-inequivalent topological saddles, i.e., the number of topological saddles in $\mathcal{M}$, is of course exactly equal to $2g-2$; in the PL case however, since we can have generic saddles with higher multiplicity, the Euler characteristic is only an upper bound for the number of topological saddles, since $\chi(\mathcal{M}) = 2-2g = \sum_{k\in\mathbb{Z}} k \cdot \#\{[p]|i(\omega,p)=k\}$.

2.3 The Close-to-Rational Case

Definition 2.8. The *irrationality degree* of a closed 1-form Ω on a piecewise smooth manifold M is the number of rationally independent integrals of Ω over any base of the homology integer 1-cycles of M: $\mathrm{irr}(\Omega) = \dim_{\mathbb{Q}} \mathrm{span}_{\mathbb{Q}}\{\int_\gamma \Omega\}$, $\gamma \in H_1(M,\mathbb{Z})$; 1-irrational forms are also called *rational*.

The structure of foliations induced on a periodic polyhedron by rational 1-forms is much simpler than in the generic case, since in this case all leaves are periodic; moreover, the appearance of a topological invariant forces all 1-forms close enough to rational to exhibit the same behavior.

Rational 1-forms are rather special because foliations induced on $\mathcal{M}$ by them are also induced by well-defined circle-valued functions on the polyhedron. Indeed, if $\{\gamma_i\}_{i=1,2,3}$ is a base for $H_1(\mathbb{T}^3,\mathbb{R})$ and N an integer big enough such that $N\int_{\gamma_i}\omega \in \mathbb{Z}$, $i=1,2,3$, and $p_0 \in \mathcal{M}$, then the function

$$f : M \longrightarrow \mathbb{S}^1,$$
$$p \mapsto \exp\left(iN\int_{p_0}^{p}\omega\right),$$

is well defined and differentiable, and its differential $df_p = iN\omega_p f(p)$ is proportional to ω. Since f is never zero, the set $df = 0$ on M coincides with the restriction to M of $\omega = 0$.

This shows that all leaves induced by rational 1-forms on a polyhedron $\mathcal{M} \subset \mathbb{T}^3$ are compact, and therefore corresponding leaves on $M \subset \mathbb{R}^3$ will be open or closed according to their homology class: Leaves homotopic to zero in $\mathbb{T}^3$ will remain so in any covering, while all others will open up in the universal covering of $\mathbb{T}^3$.

Definition 2.9. From now on we will refer to closed leaves not homotopic to zero as *periodic*, or, more generically, *open*, leaves, so that by *closed leaf* we will implicitly mean a closed leaf homotopic to zero.

Lemma 2.10. *Let ω be a rational 1-form in general position with respect to a polyhedron $\mathcal{M}$. Then no more than two open leaves can collide at any saddle point.*

Proof: The 1-form ω foliates $\mathbb{T}^3$ in a one-parameter family of embedded 2-tori, so that all open leaves at the same level of ω are parallel (i.e., they represent the same 1-cycle modulo sign); moreover, the number of open leaves on every level is even, since they are all integral and their sum must be zero.

It is easy to check, just by drawing pictures, that it is possible to have saddles of any index with closed leaves, and adding a single pair of open leaves does not change this situation; their presence, however, does not allow the presence of any other pair, since at any saddle point the two extremes of the same (critical) leaf must appear next to each other, and this is of course impossible for all open leaves, apart for the two most external ones (see Figure 2). □

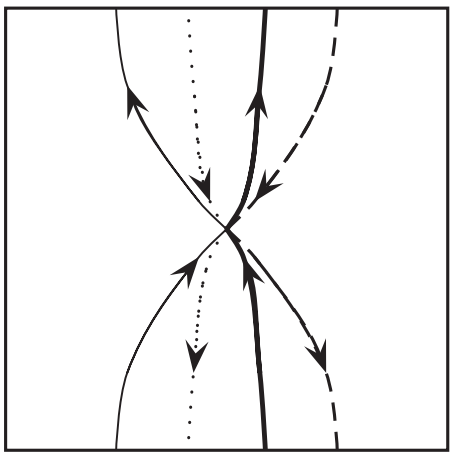

FIGURE 2. A (hypothetical) saddle point where four open leaves meet simultaneously. The two branches of each of the inner open leaves cannot be adjacent to each other, so this picture cannot come from the section of a locally Euclidean surface.

Theorem 2.11. *The Dynnikov decomposition induced on $\mathcal{M}$ by any constant rational 1-form in general position is Zorich.*

Proof: Since all (noncritical) leaves induced by ω on $\mathcal{M}$ are circles, the critical points of ω on $\mathcal{M}$ determine a subdivision of the polyhedron in the connected sum of a finite number of cobordisms. In the smooth case, the only nontrivial cobordisms are elementary (i.e., pants-like), and Zorich [Zorich 84] proved that on the boundary of each pair of pants there is at least one of the three boundary loops that is homotopic to zero (in $\mathbb{T}^3$), from which the theorem follows easily.

In the polyhedral case the cobordisms are not necessarily elementary, i.e., more than two leaves may collide at a saddle point, but lemma 2.10 grants us that even in this case no more than two loops in the cobordism may be open. This shows that the components of open leaves have genus 1, and therefore the decomposition induced on $\mathcal{M}$ by ω is Zorich. □

The presence of a genus-1 rank-2 component of open leaves is a very strong condition and leads to the appearance of a topological quantity associated with the foliation [Dynnikov 97]:

Theorem 2.12. *A Dynnikov decomposition $\mathcal{Z}$ of $\mathcal{M}$ of rank* 2 *has genus* 1 *if and only if at least one of its $\mathcal{N}_j$ embedded with rank* 2 *has genus* 1*. In this case, all $\mathcal{N}_j$ are genus-*1 *components of $\mathcal{M}$, and those embedded with rank* 2 *represent, modulo a sign, the same integral nonzero* 2*-homology class in $\mathbb{T}^3$.*

Definition 2.13. The *unsigned* nontrivial indivisible homology class $l \in \mathrm{PH}_2(\mathbb{T}^3, \mathbb{Z})$ common to all rank-2 components of $\mathcal{Z}$ is called the *soul* of a Zorich decomposition $\mathcal{Z}$.

Corollary 2.14. *Let ω be a rational* 1*-form inducing on $\mathcal{M}$ a rank-*2 *Zorich decomposition $\mathcal{Z}$. Then any* 1*-form ω' close enough to ω induces on $\mathcal{M}$ a rank-*2 *Zorich decomposition $\mathcal{Z}'$, and this decomposition is homotopic to $\mathcal{Z}$. In particular, all such decompositions share the same soul; that is, the soul is a locally constant function of the pair $(\mathcal{M}, \omega)$.*

Proof: The leaves at the boundary between the genus-1 components of periodic leaves N_i of $\mathcal{M}$ and the cylinders of closed leaves are themselves closed and therefore stable under small perturbations of ω's direction; consequently, no cylinder of closed orbits will disappear in a whole neighborhood of $\omega \in \mathbb{R}\mathrm{P}^2$. Since no two leaves of a foliation can intersect, this means that open leaves are bound to genus-1 components of $\mathcal{M}$, homotopic to the $\mathcal{Z}$ ones, even for all 1-forms close enough to ω.

Finally, since the homology class l of these rank-2 genus-1 components is integral and they change continuously, this l must be the same (modulo sign) for all them. □

The soul l of a Zorich decomposition is a fundamental invariant, since it predicts the asymptotic behavior of the open leaves. Indeed, the fact that $[\mathcal{N}_i] = \pm l$ means that in $\mathbb{R}^3$, the lift $\hat{\mathcal{N}}_i$ of $\mathcal{N}_i$ lies between a pair of parallel planes perpendicular to l (seen as a direction in $\mathbb{R}^3$), so that the lift to $\mathbb{R}^3$ of an open leaf, namely an open intersection of $\hat{\mathcal{N}}_i$ with a plane perpendicular to ω (seen as a vector in $\mathbb{R}^3$), is a curve contained in a plane strip of finite width. Dynnikov [Dynnikov 92] showed how these conditions are enough for one to conclude that the leaf is actually a finite deformation of a straight line whose direction is the axis of the strip, namely $\omega \times l$.

Since the foliation induced by ω is determined just by its direction $\boldsymbol{H}$ and the soul itself can be interpreted as a direction in $\mathbb{R}^3$, once we fix a surface M we can think of the soul application as a locally constant application $\mathrm{soul}_{\mathcal{M}} : \mathbb{R}\mathrm{P}^2 \to \mathbb{R}\mathrm{P}^2$. We will show in the next section that for a generic polyhedron $\mathcal{M}$, this map is well defined on the whole projective plane except for a set of measure zero (*ergodic directions*), and its image amounts to a finite number of points.

Definition 2.15. The nonempty level sets $\mathcal{D}_l(\mathcal{M}) = \mathrm{soul}_{\mathcal{M}}^{-1}(l)$ corresponding to nonzero values of l are called *islands* or *stability zones* of $\mathcal{M}$. The union of all stability zones $\mathcal{S}(\mathcal{M}) = \cup_l \mathcal{D}_l(\mathcal{M})$ is called the *stereographic map* (SM) of $\mathcal{M}$.

Strange as it may seem, this whole construction is the natural model for a very concrete physical phenomenon, namely the magnetoresistance in normal metals, with the periodic surface being the *Fermi surface* of a metal, the 1-form being a strong, constant, homogeneous external magnetic field, and the Fermi surface's leaves being the orbits of the momenta of the metal's quasielectrons. The presence of open leaves is detectable experimentally, and so is each stability zone (provided at least two points of it are measured [Novikov and Maltsev 98]). Experimental plots of the SM were produced in the 1960s and 1970s for about thirty metals, but only recently have the first two SM, relative to the Fermi surfaces of gold and silver, been reproduced theoretically from first principles [De Leo 04b, De Leo 05a].

2.4 The Generic (3-Irrational) Case

In the previous section we showed that the structure of the foliation in a neighborhood of rational directions is rather simple: Either all leaves are closed or their lift is strongly asymptotic to a straight line [Dynnikov 97], that is, it is contained in a finite-width plane strip and crosses it from one end to the other. Still, this is far from being enough, since the soul density of an open set does not preclude its measure from being small.

We will prove in this section that the situation above is nevertheless the generic one:

Theorem 2.16. *The set of directions in $\mathbb{R}\,\mathrm{P}^2$ inducing Dynnikov decompositions of genus greater than one on a generic polyhedron $\mathcal{M}$ has measure zero.*

To prove this theorem, Dynnikov studied the structure of the foliation induced by a 3-irrational 1-form ω on a 1-parameter family of surfaces $M_e = f^{-1}(e)$, where f is a triply periodic Morse function such that for almost all values of f no more than one critical point of ω lies on the same leaf. Here, we will repeat Dynnikov's steps, assuming f to be a generic triply periodic PL function, so that almost all of its level surfaces are embedded polyhedra satisfying the same genericity condition with respect to ω.

The following two lemmas of Dynnikov extend with no change to the polyhedral case:

Lemma 2.17. *Given a triply periodic PL function, the set of values for which a constant 1-form ω induces open leaves filling rank-2 components is a closed connected interval $[e_1(\omega), e_2(\omega)]$. The functions e and E are continuous on all of $\mathbb{R}\,\mathrm{P}^2$.*

Definition 2.18. The *Dynnikov index* w of a critical point c of $\omega = h^\alpha dp_\alpha$ with respect to an oriented polyhedron $\mathcal{M}$ is the product of the "Hamiltonian" index of the critical point ($+1$ for centers and -1 for saddles) times the sign of the scalar product between $\boldsymbol{H} = (h^\alpha)$ and any of the normals to the faces adjacent to the critical vertex.

Lemma 2.19. *The curve $\gamma_{\omega,f}(e) = \sum_i w_i c_i(e)$ is a well-defined loop in $\mathbb{T}^3$, and the quantity*

$$\overline{\chi}_{\omega,f}(e) = \int_{-\infty}^{e} i^*_{\gamma_{\omega,f}}\omega$$

is equal to the density of closed leaves on any leaf induced by ω on $\mathbb{T}^3$.

Moreover, if $\{h_j^+\}$ (respectively $\{h_k^-\}$) are the heights of the "positive" (respectively "negative") cylinders of closed leaves on $\mathcal{M}$, namely those that contain points with smaller (respectively larger) values of f, it turns out that it is possible to choose $\mathbb{R}^3$ representatives $\hat{c}_i$ of the critical points such that

$$\overline{\chi}_{\omega,f}(e) = \sum_i \langle \boldsymbol{H}, w_i \hat{c}_i(e)\rangle = \sum h_j^+ - \sum h_k^-.$$

The weighted sum of all critical points also contains the center-saddle pairs, which can be removed by homotopy and have nothing to do with the topology of the foliation. If we do not include them in the summation, we are left with a new quantity that tells us the density of nontrivial (in $\mathcal{M}$) closed leaves, and therefore we are able to determine whether there are cylinders of nonzero height. In particular, to have a full ergodic situation, i.e., a leaf dense on the whole of $\mathcal{M}$, this function must necessarily be zero.

Lemma 2.20. *Let $\{c_i^{\mathrm{top}}\}$ be the subset of the topological saddles of the pair (ω, f). Then the reduced curve $\gamma^{\mathrm{top}}_{\omega,f}(e) = \sum_i w_i c_i^{\mathrm{top}}(e)$ is also well defined in $\mathbb{T}^3$, and the reduced Euler density*

$$\overline{\chi}^{\mathrm{top}}_{\omega,f}(e) = \int_{-\infty}^{e} i^*_{\gamma^{\mathrm{top}}_{\omega,f}}\omega$$

is equal to the density of closed leaves that are not homotopic to zero in $\mathcal{M}$.

The function $\overline{\chi}^{\mathrm{top}}_{\omega,f}$ is a strictly increasing noncontinuous PL function with respect to e; its points of discontinuity are exactly the nonproper values of f.

Proof: The original curve $\gamma_{\omega,f}(e)$ is well defined because $\sum w_i = 0$ [Dynnikov 97]; since the Dynnikov indices of a saddle-center pair are opposite, the restricted sum $\sum_{\mathrm{top}} w_i$ is still zero, and therefore $\gamma^{\mathrm{top}}_{\omega,f}(e)$ is well defined.

Considering only the sum of the topological saddles is equivalent to canceling from a leaf all closed leaves that are homotopic to zero in $\mathcal{M}$, so that the averaged Euler characteristic is now relative to only the nontrivial closed leaves.

Now consider a positive cylinder. Since the values inside are smaller with respect to the values on the cylinder, the height of the cylinder increases with e; for the same reason, negative cylinders decrease their height. Since $\overline{\chi}^{\mathrm{top}}_{\omega,f}(e) = \sum_{\mathrm{top}} h_j^+ - \sum_{\mathrm{top}} h_k^-$, the function $\overline{\chi}^{\mathrm{top}}_{\omega,f}(e)$ is strictly increasing at its points of continuity. The function fails to be continuous when new pairs of topologi-

cal saddles are created or destroyed, namely, in the nonproper values of f, and it jumps exactly by the height of the cylinder(s) created or destroyed. □

Corollary 2.21. *If ω induces on $\mathcal{M}_e$ a Dynnikov decomposition of genus greater than 1, then $e_1(\omega) = e_2(\omega) = e_0$, i.e., ω induces only closed leaves at any other level.*

Corollary 2.22. *At almost all levels of f the measure of the "ergodic" directions is zero.*

We have now all the ingredients to prove Theorem 2.16:

Proof: We will discuss only the case of "full ergodicity," namely the case of directions giving rise to leaves dense on the whole polyhedron; the same line of argument extends to the lesser ergodic cases.

If ω induces on $\mathcal{M}$ fully ergodic leaves, then of course no cylinder can appear, and therefore for any Morse PL function f such that $\mathcal{M} = f^{-1}(0)$, it must happen that $\overline{\chi}^{\text{top}}_{\omega,f}(0) = 0$. Let us consider now $\overline{\chi}^{\text{top}}$ as a function of ω and e. Then the surface $X = (\overline{\chi}^{\text{top}}_f)^{-1}(0) \subset \mathbb{R}\,\mathrm{P}^2 \times \mathbb{R}$, for a generic function f, is transversal to the sections $\mathbb{R}\,\mathrm{P}^2 \times \{e\}$. Indeed, in a projective chart, say $h^z = 1$, we have that $\partial_{h^x} \overline{\chi}^{\text{top}}_{\omega,f} = \sum_i w_i (c_i^{\text{top}})^x$ and similarly for h^y, so that the points on X where the gradient is zero are exactly the points where $\gamma^{\text{top}}_{\omega,f}(e) = \sum_i w_i c_i^{\text{top}} = 0$. This condition is nongeneric, and that finally proves the claim of the theorem. □

2.5 Structure of the Stereographic Map of Surfaces and Triply Periodic Functions

The results of the previous two sections show that the SM $\mathcal{S}(\mathcal{M})$ of a generic surface is the disjoint union of a countable set of open sets $\mathcal{D}_l(\mathcal{M})$ ("islands"), each labeled by an $l \in \mathrm{PH}_2(\mathbb{T}^3, \mathbb{Z})$ immersed in a sea of directions that give rise only to closed leaves. According to our intuition of the system and the numerical experiments made to date, we conjecture that generically, the number of islands is finite, but no rigorous proof of this fact exists. As a matter of fact, this structure is exactly the one guessed, from symmetry considerations, by the physicist I. M. Lifschitz and his Kharkov school about fifty years ago [Lifschitz and Peschanskii 59, Lifschitz and Peschanskii 60].

The boundaries of the islands are reached when the last pair of genus-1 components collide because of a cylinder collapse and therefore are characterized by the presence of (at least) a pair of inequivalent critical points on the same leaf. The set of these directions is the countable union of the curves $\langle \boldsymbol{H}, \hat{c}_i - \hat{c}_j \rangle = 0$, $i \neq j$; in the polyhedral case all such curves are straight lines, so that every island is actually a (not necessarily convex) polygon. There are reasons to believe that these polygons are convex for low genus, i.e., at least for genera 3 and 4, but it is easy to build examples of high-genus polyhedra with islands that are either not connected, or connected but not convex, or even connected but not simply connected.

A crucial observation by Dynnikov [Dynnikov 97] allows us also to associate an SM with triply periodic functions:

Theorem 2.23. *Let f be a Morse triply periodic function. Then if ω induces on $\mathcal{M}_{e0}$ a Zorich decomposition $\mathcal{Z}_{e0}$, it induces a Zorich decomposition $\mathcal{Z}_e$ for all $e \in (e_1(\omega), e_2(\omega))$, and all these decompositions share the same soul.*

Definition 2.24. The island $\mathcal{D}_l(f)$ corresponding to the label l is the union of the corresponding islands of the level sets of f: $\mathcal{D}_l(f) = \cup_{e \in \mathbb{R}} \mathcal{D}_l(\mathcal{M}_e)$. The SM of f is the union of all its islands: $\mathcal{S}(f) = \cup_l \mathcal{D}_l(f)$.

The SMs corresponding to functions are generically dramatically different from those corresponding to surfaces.

First of all, since rational directions necessarily induce Zorich decompositions, the set of islands $\mathcal{S}(f)$ is now always dense in $\mathbb{R}\,\mathrm{P}^2$, and of course Lemma 2.17 also tells us that in this case the sea of directions giving rise to closed leaves has dried up, since every direction either belongs to an island (or its boundary) or is ergodic.

Moreover, the following property shows that the islands can be sorted in a rather complex way:

Theorem 2.25. *Generically, every two zones meet transversally and in a countable number of points.*

Proof: No point belonging at the same time to two different zones $\mathcal{D}_{l_1}$ and $\mathcal{D}_{l_2}$ can have irrationality degree greater than 2, since it must contain the integer direction $l_1 \times l_2$. In the smooth case this would be enough, since the boundaries are smooth curves and they generically contain only a countable number of 2-irrational points and no rational points. In the PL case, however, the boundaries are actually segments of straight lines, and therefore they are actually contained in the set of the directions with irrationality degree less than one.

Nevertheless, the theorem holds in this case as well for the following reason: Since every direction at the boundary between two zones is perpendicular to the direction $l_1 \times l_2$, their set is the straight line (in $\mathbb{R}\,\mathrm{P}^2$) passing through l_1 and l_2. Generically, none of the two labels falls on the boundary, and therefore two zones can meet in a number of points not greater than the number of sides of the island with the smaller number of sides. □

Corollary 2.26. *Either there is a single zone, i.e., there exists a label l such that $\mathcal{D}_l(f) = \mathbb{R}\,\mathrm{P}^2$, or there are countably many zones and they are dense in the whole projective plane.*

Since the islands meet transversally, the nontrivial SMs will look like 2-dimensional Cantor sets. Analyzing such fractals numerically is not trivial, since it involves, in general, examining the system at several values of f, but there is a class of interesting (nongeneric) cases in which it is actually enough to analyze a single level surface of f to obtain the entire fractal picture:

Theorem 2.27. *Let $\mathcal{M}$ be a polyhedron whose interior is equal to its exterior modulo the group $G \simeq \mathbb{R}^3 \times \mathbb{Z}_2$ of translations and inversion of the three axes. Then $\mathcal{S}(\mathcal{M}) = \mathcal{S}(f)$ for any function having $\mathcal{M}$ as (connected component of a) level set.*

Proof: By symmetry, we can build a function f such that $\mathcal{M} = f^{-1}(0)$ and such that $f^{-1}(e)$ is equal to $f^{-1}(-e)$ modulo G. Since the bundles of parallel planes are also invariant by G, it turns out that for such an f the interval $I(\omega)$ of existence of open orbits relative to any 1-form ω is of the form $I(\omega) = [-a(\omega), a(\omega)]$, and therefore $0 \in I(\omega)$, $\forall \omega \in \mathbb{R}\,\mathrm{P}^2$, i.e., at the zero level every ω induces open leaves. □

In particular, all triply periodic functions f such that $f(c-x, c-y, c-z) = -f(x, y, z)$ belong to this class, and indeed the only fractals analyzed numerically to date are relative to this kind of function.

Finally, we cite an important property that ties the set of all labels relative to the islands of an SM of a function with the set of ergodic directions [De Leo 04a, De Leo 05b]:

Theorem 2.28. *The closure of the set of all labels is the disjoint union of the set of all zone boundaries and the set of ergodic directions.*

3. A CONCRETE CASE STUDY

As pointed out in the previous section, to date, a picture of the fractal has been numerically produced for only two functions: an analytical one and a piecewise quadratic one [De Leo 04a].

The analytical function is $\mathfrak{f}(x, y, z) = \cos(2\pi x) + \cos(2\pi y) + \cos(2\pi z)$, invariant with respect to translations by integers $\Gamma = \mathbb{Z}^3 \subset \mathbb{R}^3$, which gives rise to genus-3 level surfaces in the range $(-1, 1)$ and spheres at every other noncritical level. This function represents the simplest nontrivial case possible from the topological point of view, since any triply periodic connected surface of genus less than 3 lies between two parallel planes, and therefore the aymptotics of plane sections are easily found. Its zero level is rather special: It is known as the Schwarz primitive function (or plumber's nightmare) and was studied by Schwarz in 1890 as one of the first examples of a triply periodic minimal surface.

From the computational point of view, $\mathcal{SP} = \mathfrak{f}^{-1}(0)$ has three important properties:

SP1. Its interior is a translate of its exterior, so that $\mathcal{S}(\mathcal{SP}) = \mathcal{S}(\mathfrak{f})$.

SP2. It is invariant with respect to the natural action of the tetrahedral group T_d on the unitary cube, so that the whole SM can be obtained, for example, by extending by symmetry to all of $\mathbb{R}\,\mathrm{P}^2$ the data obtained for the triangle with vertices $[(0, 0, 1)]$, $[(1, 0, 1)]$, and $[(1, 1, 1)]$.

SP3. The two cylinders, one negative and one positive, have the same height, so that it is enough to examine just one of the four topological critical points at the base of the two cylinders to retrieve all information about the structure of the foliation.

The piecewise quadratic function is $\mathfrak{g}(x, y, z) = \sum \bar{\mathfrak{g}}(x^i)$, where $\bar{\mathfrak{g}}$ is the simplest piecewise quadratic function having the same symmetries as those of the cosine function. Its level sets have the same behavior as the function above, but the expression of the critical points as a function of the direction of the 1-form and the level of the function is so simple that a comparison between the analytical and numerical data is possible also at levels different from zero.

Nevertheless, in both cases the number of triangles needed to describe the surface in sufficient detail is so large (between 10^5 and 10^6) that it is impossible to improve the resolution of the results obtained in [De Leo 04a], at least until a new algorithm is found or some ad hoc trick is discovered.

3.1 The Polyhedron

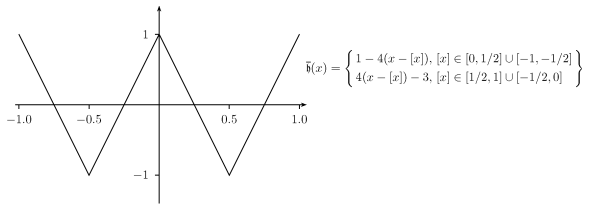

FIGURE 3. The simplest PL approximation of the cosine function.

A natural way to improve the resolution of the numerical analysis of the problem is to consider PL functions, since in this case the number of triangles needed to describe their level surfaces can be as low as order 10^1. Moreover, the description of the surfaces is in this case exact rather than an approximation.

The simplest case to study, in order to take advantage of the extremely convenient properties evidenced in the $\mathcal{SP}$ case, is the PL function $\mathfrak{h}(x,y,z) = \sum_i \bar{\mathfrak{h}}(x^i)$, where $\bar{\mathfrak{h}}$ is the function shown in Figure 3. The polyhedron $\mathcal{P}_0 = \mathfrak{h}^{-1}(0)$ is a PL embedding of a genus-3 surface in $\mathbb{T}^3$ having 8 hexagonal faces, 20 edges and 12 vertices; the smallest triangulation for $\mathcal{P}_0$ takes $4\cdot 8 = 32$ triangles, $32\cdot 3/2 = 48$ edges, and 12 vertices, which gives the expected Euler characteristic $\chi(\mathcal{P}_0) = F-E+V = -4 = 2-2\cdot 3$. The basic cell of the lift $\overline{\mathcal{P}}_0 \subset \mathbb{R}^3$ in the unit cube is a truncated octahedron (Figure 4); it is noteworthy to notice that like its smooth analogue, this surface is, in the discrete sense, also a minimal surface [Wayne 05].

Since exactly four edges meet at every vertex, given any 1-form ω in general position with respect to $\mathcal{P}_0$, only saddles with index -1 can arise, and therefore there are always four vertices that are critical for such ω. If the point p_1 is such a vertex, then the remaining three critical points are $p_{2,3} = \left(\frac{1}{2},\frac{1}{2},\frac{1}{2}\right) \pm p_1$ and $p_4 = (1,1,1) - p_1$. In particular, in our numerical study we sample the set of 1-forms $\omega = (h^x, h^y, h^z)$ such that $h^x/h^z \in [0,1]$ and $h^y/h^z \in [0,1]$, for which the four critical points are $(0,.5,.75)$, $(.5,1,1.25)$, $(.5,0,.25)$, and $(1,.5,.25)$.

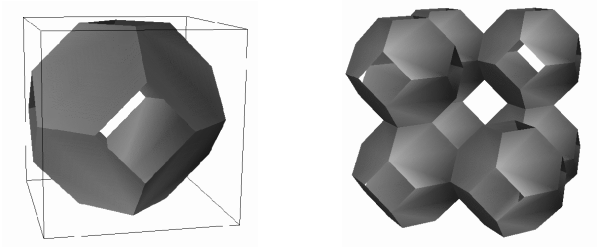

FIGURE 4. Plot of the truncated octahedron inside $\mathbb{T}^3$ (left) and of part of its image in the universal covering.

3.2 The Algorithm

In order to generate an approximate picture of a fractal, it is enough to produce an algorithm able to evaluate the label, if any, associated with a given 1-form. Since obviously no calculator can deal with irrational numbers, the numerical study will be limited to rational 1-forms; fortunately, this is not a big restriction, since in any case, rational directions are dense in every stability zone.

Note that the algorithm used for the numerical study of the PL case is a simplified version of the more general algorithm we developed to study smooth surfaces of genus three [De Leo 04a], since in this case we know a priori the position of all critical points and moreover, we know their position exactly, so that we do not have to correct "by hand" the topology of the critical section.

The basic idea to retrieve the label, as suggested to me by I. Dynnikov, is that the soul $l \in \mathrm{PH}_2(\mathbb{T}^3,\mathbb{Z})$ associated with the Zorich decomposition $\mathcal{Z}$ induced by $\omega \in \mathcal{D}_l$ is in one-to-one correspondence with the rank-2 sublattice of $H_1(\mathbb{T}^3,\mathbb{Z})$ obtained as the image, through the map $i_* : H_1(\mathcal{M},\mathbb{Z}) \to H_1(\mathbb{T}^3,\mathbb{Z})$, of the open leaves in $\mathcal{M}$ that have zero intersection number with the closed leaves pop-

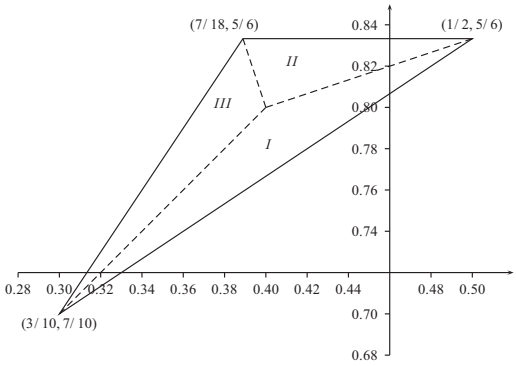

FIGURE 5. A close-up of the island $\mathcal{D}_{(2,4,5)}(\mathcal{P}_0) = \mathcal{D}_{(2,4,5)}(\mathfrak{h})$. Inside the island the pairs of critical points at the base of each cylinder are locally constant. In correspondence to each side, there are three different pairings sorted in open subsets, labeled in the picture by roman numerals, separated by straight line segments corresponding to directions ω for which the bases of the positive and negative cylinders collide, resulting in a saddle connection between two critical points. These three segments meet in the single point $(.4,.8)$, which in this case happens to be exactly the direction of the label (this property, though, is not generic). In Figure 6 we show in detail the transition between two different pairs within this island.

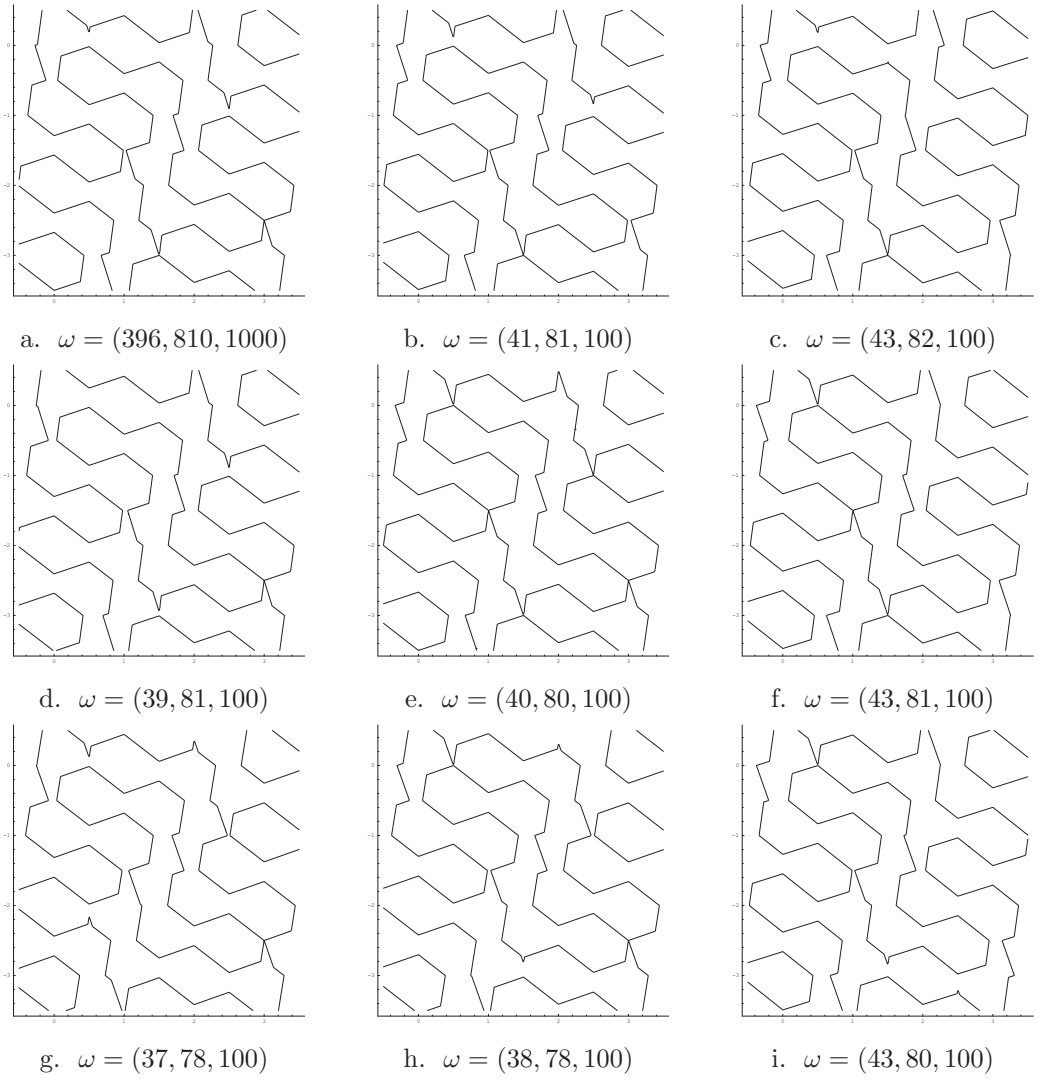

a. $\omega = (396, 810, 1000)$ b. $\omega = (41, 81, 100)$ c. $\omega = (43, 82, 100)$

d. $\omega = (39, 81, 100)$ e. $\omega = (40, 80, 100)$ f. $\omega = (43, 81, 100)$

g. $\omega = (37, 78, 100)$ h. $\omega = (38, 78, 100)$ i. $\omega = (43, 80, 100)$

FIGURE 6. Significant examples of critical sections of $\mathcal{P}_0$ for $\omega \in \mathcal{D}_{(2,4,5)}(\mathcal{P}_0)$. At the "center" of the island (e) there is a saddle connection between the four critical points (starting from the highest and turning clockwise) $N = p_1$, $E = p_2 + (3, -3, 1)$, $S = p_1 + (1, -3, 2)$, and $W = p_2 + (1, -2, 1)$. In the subzones I, II, and III (Figure 5) the pairs of critical points at the base of the positive cylinder are, respectively, N and $p_4 + (2, -4, 1)$ (i), S and $p_3 + (2, -2, 0)$ (b,c), and E and $p_3 + (-1, -1, 0)$, (d, g). The separating segments correspond respectively to saddle connections between the pairs of critical points S and W (a), N and S (f), and N and E (h).

ulating the cylinders of $\mathcal{Z}$. Indeed, every cycle lying on the interior of an N_j component of $\mathcal{Z}$ has no intersection with the closed leaves that form the cylinders $\mathcal{C}_i$, and since all N_j are homologous to each other (modulo sign), the image of all these cycles in $\mathbb{T}^3$ must have rank 2; on the other hand, there is an obvious one-to-one correspondence between rank-2 sublattices of $H_1(\mathbb{T}^3, \mathbb{Z}) = \Gamma \simeq \mathbb{Z}^3$ and the 2-tori embedded in $\mathbb{T}^3 = \mathbb{R}^3/\Gamma$, since every such 2-torus can be spanned by a pair of independent rational directions and vice versa.

The following algorithm works for genus-3 polyhedra satisfying properties SP1–SP3, in particular for $\mathcal{P}_0$:

Algorithm 3.1.

Input: $\mathcal{M}$, the polyhedron; $\omega = (l, m, n) \in \mathbb{Z}^3$, the 1-form; x, a critical point of ω with respect to $\mathcal{M}$; $\pi_{\omega,x}$, the plane perpendicular to (l, m, n) and passing through x.

Output: $c_{1,2}$, the two critical loops;[2] $h_{1,2}$, the homology classes of $c_{1,2}$ in $\mathcal{M}$; $H_{1,2}$, the homology classes of $c_{1,2}$ in $\mathbb{T}^3$.

[2]Since numerically we can study only the 1-rational case, in $\mathbb{T}^3$ the saddles are always wedges of circles, i.e., all critical branches close back to the critical point; according to whether these loops are or are not homotopic to zero in $\mathbb{T}^3$, their $\mathbb{R}^3$ lift will be open or closed.

N1. Retrieve the intersection between M and $\pi_{\omega,x}$.

N2. Check that there are exactly four critical branches and follow them by periodicity; otherwise exit.

N3. If no other critical point is met along the path, so that the four branches are arranged in a pair of critical loops, store the two loops in the variables $c_{1,2}$; otherwise exit.

N4. Evaluate the homology class of $c_{1,2}$ in $\mathbb{T}^3$ and in M (this is actually done while executing step N2 to speed up the computation time).

N5. If the saddle is half-open, i.e., if exactly one among $H_{1,2}$ is zero, then associate with ω the complementary h triple; otherwise exit.

The main outcome of the algorithm is of course the label associated with ω. The fact that this label is a triple of integers is very important, since an integer evaluated numerically with an error smaller than 0.5 becomes actually an exact measure.

We implemented this algorithm in a C++ library named NTC.[3] built over an Open Source C++ library named VTK.[4] The choice of the language comes from the fact that VTK provides the basic geometric environment and algorithms needed by the problem, mainly the capability of generating meshes for isosurfaces and evaluating intersections between geometric objects. The inheritance mechanism of the C++ language allows one to use transparently all functions of a library; hence we used VTK as a starting point and implemented in NTC the routines to deal with periodicity and evaluate the homology classes.

No serious attempt to evaluate the error on such calculations has been made to date, since no need for it has arisen. Indirect evidence was used to check the reliability of the results:

- the agreement of the largest zones with their analytical boundary (Figure 9) obtained through the independent Algorithm 3.2;
- the symmetry of the final picture with respect to the diagonal (Figure 11), symmetry that was in no way used in the numerical calculations;
- the agreement of the fractal picture with the plot of the labels (Figure 15).

[3]Available at http://ntc.sf.net/
[4]Available at http://www.vtk.org/.

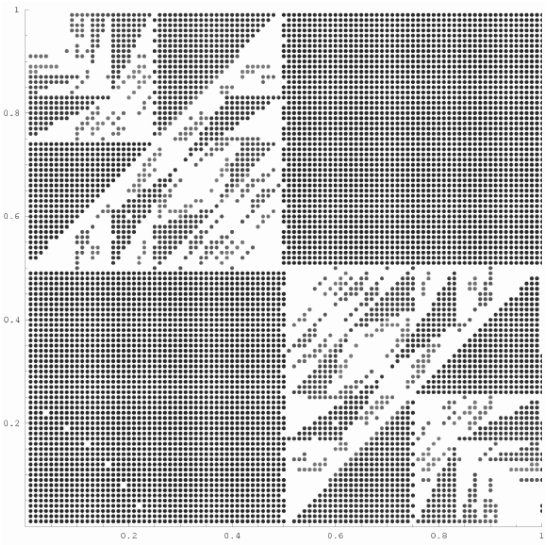

FIGURE 7. Numerical plot of the square $[0,1]^2$ of the SM $\mathcal{S}(\mathcal{P}_0) = \mathcal{S}(\mathfrak{h})$ in the projective chart $h^z = 1$ at a resolution $r = 10^2$. The color of the islands goes from blue to red as the norm of the label grows. In the picture are displayed the 106 islands with at least four points out of the total of 1741 islands found. The missing points that are possible to see in the interior of some of the islands are due to a failure of the numerical Algorithm 3.1, e.g., because of the presence of saddle connections. The running time for the 10^4-step cycle needed to retrieve these data takes about 1 h on a Pentium $\approx$ 1-GHz CPU.

The exploration of the SM was performed in the square $[0,1]^2$ of the projective chart $(\omega_x/\omega_z, \omega_y/\omega_z)$ by evaluating the label associated with every direction at the vertex of a uniform grid of step size r, and this was repeated for the values $r = 10^2, 10^3, 10^4$. Samplings with $r = 10^2, 10^3$ have been successfully performed as well for the previous two functions [De Leo 04a], but the CPU time needed to reach $r = 10^4$ in that case was way too long. It is because of the rather small number of triangles needed to describe $\mathcal{P}_0$ that the computation became feasible.

3.3 Numerical Results for $r = 10^2$

This resolution is the lowest that allows us to obtain a hint of the structure of the fractal. About r sections are needed to follow the critical branches for a generic direction (m, n, r), $m, n \in \{1, \ldots, r\}$, which takes a time of 0.5 s on an $\approx$ 1-GHz CPU for the evaluation of a single label and $10^4 \times 0.5$ s ≈ 1 h for sampling the 10^4 directions of the grid (Figure 7).

Even from this rough picture a further symmetry of the picture is rather evident, namely that with respect to the antidiagonal of the square. This symmetry does

not come, like the others, from the tetrahedral group T_d; it is rather of a topological nature. The numerical evidence is that if a 1-form (m, n, r) is labeled by L, then its symmetric form $(1-n, 1-m, r)$ is labeled by $L + (1, 1, 0)$, but no proof of this fact is known.

In order to verify the correctness of the algorithm, we found "by hand" the analytical boundaries for the largest zones and compared them with the numerical results (Figure 9). The following algorithm, aimed at the cases similar to the cosine one, is a slight modification of the original algorithm introduced and used by Dynnikov in [Dynnikov 96]:

Algorithm 3.2.

A1. Fix a 1-form $\omega = (m, n, r) \in \mathbb{Z}^3$ inside some zone (e.g., extracting it from the experiment, guessing it from symmetry arguments, or simply by trial and error).

A2. Retrieve the critical section of ω passing through one of its critical points c and make sure that it is half open; otherwise exit.

A3. Evaluate the homology class l of the closed critical leaf C.

A4. Rotate ω around some direction until the cylinder of which C is a base collapses and identify the critical point c' that is now connected to c through a saddle connection.

A5. The equation $\langle \boldsymbol{H}, c - c' \rangle = 0$ contains one of the sides of the island; follow it in one direction until four critical points fall over the critical closed leaf. This is the point at which a side and a new one meet; repeat this step until the island boundaries close up on themselves.

Since the genus is three, only two cylinders may appear, and they will be of opposite sign. The pairs of critical points at the base of cylinders are locally constant; in the square under investigation the pairings are p_1, p_4 for the positive cylinder and p_2, p_3 for the negative one, so that boundaries are always given by an equation like $\langle \boldsymbol{H}, p_1 - p_4 + L \rangle = 0$. See Figures 5 and 6 for a concrete example worked out in detail.

Finally, a picture of the whole fractal can be obtained through the natural free action on $\mathbb{R}\,\mathrm{P}^2$ of the tetrahedral group T_d, whose order is 24 (Figure 6(d)).

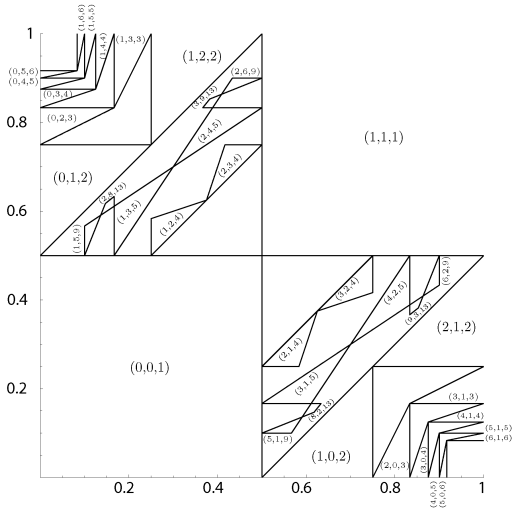

FIGURE 8. Analytical boundaries of the largest islands found in Figure 7 obtained using Algorithm 3.2. All of their boundaries are straight line segments, and the corresponding label has been reported, when possible, inside the zone itself.

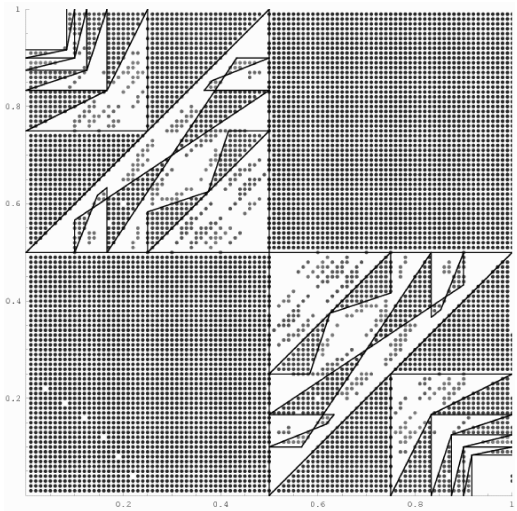

FIGURE 9. Comparison between analytical (Figure 8) and numerical (Figure 7) boundaries.

3.4 Numerical Results for $r = 10^3$

This is the highest resolution reached in [De Leo 04a]. In this case, for each generic direction are needed about 10^3 sections to follow the critical branches, which takes a time of about 2.5 s on an $\approx$ 1-GHz CPU for the evaluation of a single label and therefore about $5 \times 10^2 \times 1$ h ≈ 1 month for sampling the 10^6 directions of the grid (Figure 11).

A time on the order of a month is of course rather long, but thanks to the diffusion of the Linux OS, and therefore of the possibility to build big Unix clusters in-

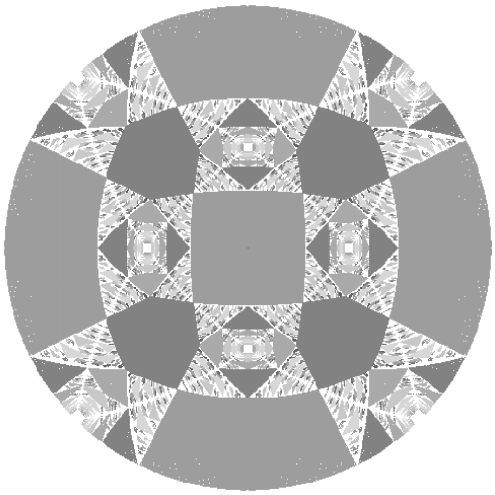

FIGURE 10. Fractal image obtained by letting the tetrahedral group T_d act on the square in Figure 7 and then projecting on the disk through the stereographic map.

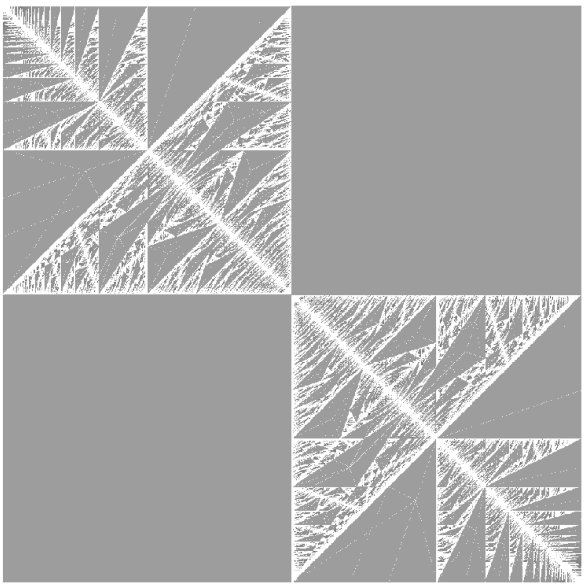

FIGURE 11. Numerical plot of the square $[0,1]^2$ of the SM $\mathcal{S}(\mathcal{P}_0) = \mathcal{S}(\mathfrak{h})$ in the projective chart $h^z = 1$ at a resolution $r = 10^3$. The color of the islands goes from green to red as the norm of the label grows. In the picture are displayed the 1625 islands with at least four points out of the total of 10 725 islands found. The missing points that are possible to see in the interior of some of the islands are due to a failure of the numerical Algorithm 3.1, e.g., because of the presence of saddle connections. The running time for the 10^6-step cycle needed to retrieve these data takes about one week on a Pentium 1-GHz CPU.

expensively using PCs rather than workstations, this is still not too bad, since it is easy to lower by a factor of 10 the computation time just by dividing the cycle over many PCs. In this way, the running time goes down to just three days, which is a rather acceptable time.

We point out that the situation is radically different in the smooth case. Indeed, in that case there is another variable to consider, namely, the resolution of the mesh of the surface, which must be increased together with the grid resolution to avoid errors in the topology of the curve. The plane sections giving the complete intersection between the surface and a 2-torus with homology class (l, m, r), $l, m \in [0, r]$, can be as close as $1/r$, and therefore, if the mesh is too rough, there is the risk that the program will jump onto the wrong slice.

Concrete tests show that a mesh resolution of 30, meaning that the mesh is produced by dividing the unit cube with a 30^3 uniform grid, is enough for the $r = 10^2$ case, but it must be raised to at least 60 for the $r = 10^3$ case, increasing the time for a single label evaluation to about 15 s, an order of magnitude greater than in the PL case. This brings the time back to about 3 months for the execution, which is indeed the order of the time spent for the $r = 10^3$ calculations made for S[De Leo 04a].

From Figure 11 the symmetry with respect to the antidiagonal is rather evident. Apart from this, the picture looks qualitatively very similar to the pictures found in the previous two cases at the same resolution.

3.5 Numerical Results for $r = 10^4$

Increasing the resolution by another order of magnitude, we increase by an order of magnitude the number of sections needed to follow a generic leaf, resulting in another factor of 5 in the running time for a single evaluation of a 1-form label, which is now about 10 s, so that the total running time on a 1-GHz CPU reaches $5 \cdot 10^2 \cdot 20$ d $= 10^4$ d ≈ 30 years.

Such a long running time is rather intimidating and suggests that there is no hope to go up by an order of magnitude in resolution without changing some significant algorithmic step. Nevertheless, this long time can be once again brought down to something reasonable by running the code on 20 PCs and by restricting the numerical analysis to the upper triangle of the square $[0, .5] \times [.5, 1]$ (which reduces computational time by a further factor of 8). Thanks to all these measures, the running time goes down by two orders of magnitude, reaching about three months, which is indeed about the time that took us to collect the $r = 10^4$ data.

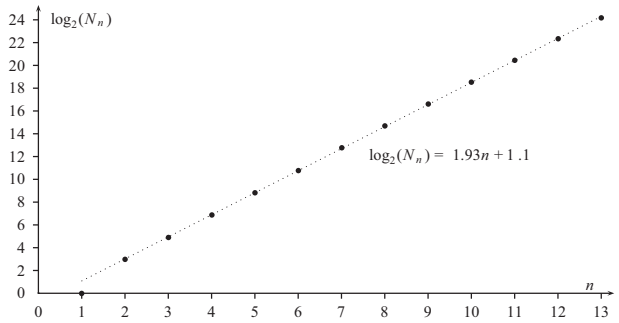

FIGURE 12. Box-counting evaluation of the Hausdorff dimension of $\mathcal{S}(\mathcal{P}_0)$ with the $r = 10^4$ data. Here N_n is the number of squares of side 2^n needed to cover the complement of $\mathcal{S}(\mathcal{P}_0)$; the angular coefficient of the linear fit provides the dimension estimate. The same evaluation made with the $r = 10^3$ data gives a very accurate result, and also restricting the fit by canceling a few points at the extremes does not change the estimate by much.

Note that for the smooth cases this would not be enough, since we must raise the mesh resolution to 10^2, and the time for the single evaluation then goes up to 50 s; this, together with the fact that there is a further factor of 2 due to the lack of symmetry with respect to the antidiagonal, brings the total running time to about 3 years, definitely not realistically affordable.

In Figures 16–30, contained in the online appendix [De Leo 06] and in the arXiv version of this paper [De Leo 05c], are shown the numerical results, from which it is clear beyond any doubt that the SM has a fractal-like self-repeating structure, even though no explicit construction is known for it.

Numerical evaluations of the Hausdorff dimension $d_{\mathcal{P}_0}$ of the fractal set, namely the complement of $\mathcal{S}(\mathcal{P}_0)$ in $\mathbb{R}\,\mathrm{P}^2$, have been performed using the box-counting technique (Figure 12) and a sort of "area distribution" technique (Figure 13).

The first, and most reliable, technique involves partitioning the square $[0,1]^2$ into 2^n identical squares and evaluating the number of squares needed to cover the fractal. With the $r = 10^3$ data, n can get as large as $\log_2(10^3) \simeq 10$, and a linear fit gives an evaluation of $d_{\mathcal{P}_0} \approx 1.93$. The $r = 10^4$ data allow n to go up to 13, representing deeper results on the Hausdorff dimension of such fractals to date. Figure 12 shows that the scaling law is linear to a high degree of accuracy, and a linear fit gives again an estimate of $d_{\mathcal{P}_0} \approx 1.93$, making us rather confident in the accuracy of this numerical result.

The second technique involves counting the number of zones whose size is between b^n and b^{n+1}, where $b > 1$.

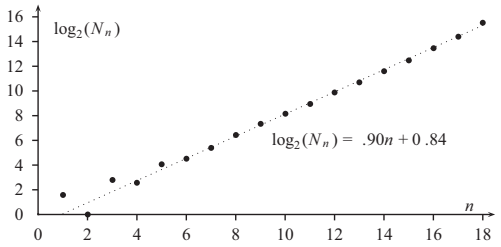

FIGURE 13. "Area distribution" evaluation of the Hausdorff dimension of $\mathcal{S}(\mathcal{P}_0)$ with the $r = 10^4$ data. Here N_n is the number of islands whose area is between 2^{-n} and 2^{-n-1}. The estimate of the Hausdorff dimension is given by the double of the angular coefficient of the linear fit and varies by ± 0.2 by restricting even by a little the number of points on which the fit is made.

Computations with several small values, between 2 and 1, were performed, and all of them give a rough estimate of $d_{\mathcal{P}_0} \approx 1.76$, further than expected from the evaluation given by box counting. The reason for this disagreement is not clear to us. It did not manifest itself in the smooth cases [De Leo 03], but it is not impossible that it may be due simply to the fact that this evaluation stabilizes at higher resolutions and therefore is at the current state of things more or less unreliable (as simple numerical tests testify).

Finally, the $r = 10^4$ data provide enough detail to test numerically a conjecture by Dynnikov [Dynnikov 99] and give graphical evidence of Theorem 2.28. The Dynnikov conjecture claims that the areas of the islands satisfy a

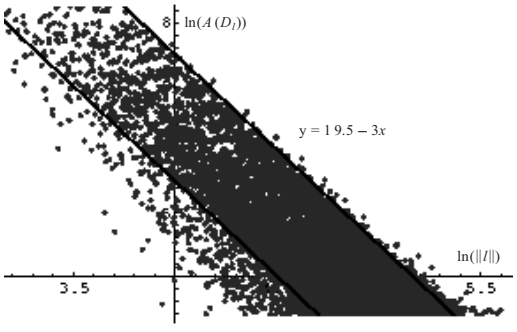

FIGURE 14. Logarithmic plot of the islands' areas versus the norm of the corresponding label. There is a very good agreement between the numerical data and the Dynnikov conjecture claiming that $\mathcal{A}(\mathcal{D}_l) \leq C/\|l\|^3$.

relation $\mathcal{A}(\mathcal{D}_l(\mathcal{M})) \leq C/\|l\|^3$ for some positive real number C depending only on the polyhedron $\mathcal{M}$; the numerical data suggest that the exponent 3 cannot be improved any further (see Figure 14).

According to Theorem 2.28, the closure of the set of labels of $\mathcal{S}(\mathcal{P}_0)$, seen as points of $\mathbb{R}\,\mathrm{P}^2$, is equal to the complement of the interior of the islands. At this resolution about $5 \cdot 10^5$ different labels are found, and their image in $\mathbb{R}\,\mathrm{P}^2$ (Figure 15) is one of the best indirect checks of the correctness of the library NTC.

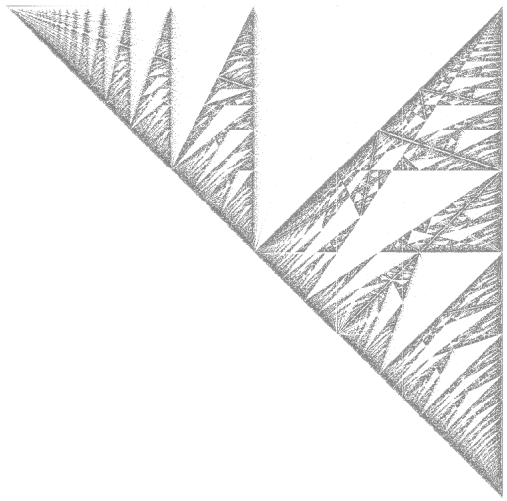

FIGURE 15. Plot of the labels of the 494 041 islands found in the numerical analysis, at the resolution $r = 10^4$, of the portion of the SM $\mathcal{S}(\mathcal{P}_0)$ contained in the triangle of vertices $(0, 1)$, $(.5, 1)$, and $(.5, .5)$ in the projective chart $h^z = 1$. According to Theorem 2.28, the closure of the set of labels is equal to the complement of the interiors of the islands; the striking closeness of the two pictures is one of the best indirect tests of the correctness of the implementation of Algorithm 3.1 in our C++ library NTC.

4. CONCLUSIONS

We have proved in this paper that the structure of foliations induced on triply periodic embedded polyhedra by constant 1-forms is identical to that induced on smooth surfaces, and we believe that the same line of argument can also be used to further generalize the theorems to embedded piecewise smooth surfaces. This fact extends to "Morse PL functions" the association of an SM; in the most interesting cases, such SM have a fractal nature (this condition is true for an open subset of all triply periodic PL functions, e.g., for all triply periodic functions close enough to $\mathfrak{h}$).

Surfaces satisfying property SP1 (see Section 3) are particularly rich, their SM being equal to the SM of any function having them as a level set. We exploited this fact by studying the case of the triply periodic surface $\mathcal{P}_0$ obtained by extending a truncated octahedron in the three coordinate directions. The simplicity of the triangulation of the surface allowed us to improve by an order of magnitude with respect to the results obtained in [De Leo 04a] the resolution of the numerical analysis of the fractal and, as a consequence, to perform several numerical tests on conjectures and theorems.

ACKNOWLEDGMENTS

We wish to thank first of all S. P. Novikov for introducing the subject. We are also deeply indebted to S. P. Novikov and I. A. Dynnikov for their interest in our work and for several precious and fruitful scientific suggestions and discussions that were essential for the present work. In particular, we thank I. A. Dynnikov for suggesting the analytical algorithm to retrieve the islands' boundaries. All numerical calculations were made with computers kindly provided by the INFN section of Cagliari (www.ca.infn.it) and CRS4 (www.crs4.it). We also acknowledge financial support from the Cagliari section of INFN and from the Departments of Physics and Mathematics of the University of Cagliari.

REFERENCES

[Banchoff 67] T. F. Banchoff. "Critical Points and Curvature for Embedded Polyhedra." *J. of Differential Geometry* 1 (1967), 245–256.

[Banchoff 70] T. F. Banchoff. "Critical Points and Curvature for Embedded Polyhedral Surfaces." *Amer. Math. Monthly* 77 (1970), 475–485.

[De Leo 03] R. De Leo. "Numerical Analysis of the Novikov Problem of a Normal Metal in a Strong Magnetic Field." *SIADS* 2:4 (2003), 517–545.

[De Leo 04a] R. De Leo. "Characterization of the Set of 'Ergodic Directions' in the Novikov's Problem of Quasi-electrons Orbits in Normal Metals." *RMS* 58 (2004), 1042–1043.

[De Leo 04b] R. De Leo. "Topological Effects in the Magnetoresistance of Au and Ag." *Physics Letters A* 332 (2004), 469–474.

[De Leo 05a] R. De Leo. "Proof of Dynnikov's Conjecture on the Location of Stability Zones in the Novikov Problem on Planar Sections of Periodic Surfaces." *RMS* 60:3 (2005), 566–567.

[De Leo 05b] R. De Leo. "First-Principles Generation of Stereographic Maps for High-Field Magnetoresistance in Normal Metals: An Application to Au and Ag." *Physica B* 362 (2005), 62–75.

[De Leo 05c] R. De Leo. "Topology of Plane Sections of Periodic Polyhedra with an Application to the Truncated Octahedron." math.DG/0502219, 2005.

[De Leo 06] R. De Leo. "Appendix to This Paper Containing High-Resolution Images." Available online at http://www.expmath.org/expmath/volumes/15/15.1/deleo/pol-appendix.pdf, 2006.

[Dynnikov 92] I. V. Dynnikov. "Proof of the S. P. Novikov Conjecture on the Semiclassical Motion of an Electron." *Usp. Mat. Nauk (RMS)* 57:3 (1992), 172.

[Dynnikov 96] I. V. Dynnikov. "Surfaces in 3-Torus: Geometry of Plane Sections." In *Proc. of ECM2, BuDA*, 1996.

[Dynnikov 97] I. V. Dynnikov. "Semiclassical Motion of the Electron. A Proof of the Novikov Conjecture in General Position and Counterexamples." In *Amer. Math. Soc. Transl. Ser. 2*, 179, pp. 45–73. Providence, RI: AMS, 1997.

[Dynnikov 99] I. V. Dynnikov. "The Geometry of Stability Regions in Novikov's Problem on the Semiclassical Motion of an Electron." *RMS* 54:1 (1999), 21–60.

[Lifschitz and Peschanskii 59] I. M. Lifschitz and V. G. Peschanskii. "Metals with Open Fermi Surfaces I." *JETP* 8 (1959), 875.

[Lifschitz and Peschanskii 60] I. M. Lifschitz and V. G. Peschanskii. "Metals with Open Fermi Surfaces II." *JETP* 11 (1960), 137.

[Novikov 82] S. P. Novikov. "Hamiltonian Formalism and a Multivalued Analogue of Morse Theory." *Usp. Mat. Nauk (RMS)* 37:5 (1982), 3–49.

[Novikov and Maltsev 98] S. P. Novikov and A. Ya. Maltsev. "Topological Phenomena in Normal Metals." *Usp. Fiz. Nauk* 41:3 (1998), 231–239.

[Wayne 05] R. Wayne. "Infinite Periodic Discrete Minimal Surfaces without Self-Intersections." *BJGA* 10:2 (2005), 106–128.

[Zorich 84] A. V. Zorich. "A Problem of Novikov on the Semiclassical Motion of Electrons in a Uniform Almost Rational Magnetic Field." *Usp. Mat. Nauk (RMS)* 39: 5 (1984), 235–236.

Roberto De Leo, Department of Physics and INFN, University of Cagliari, Cagliari, Italy (roberto.deleo@ca.infn.it)

Received February 11, 2005; accepted July 22, 2005.