

Visualization of Topological Structures in Area-Preserving Maps

Xavier Tricoche, *Member, IEEE*, Christoph Garth, *Member, IEEE*, and Allen Sanderson

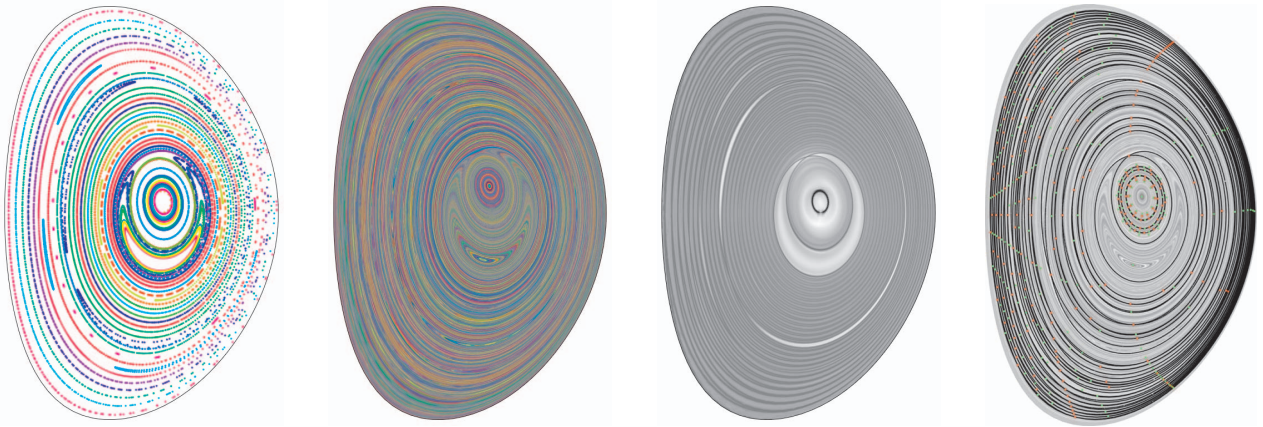


Fig. 1. Topological structure of the Poincaré section of the magnetic field on the poloidal plane of a Tokamak fusion reactor. From left to right, the structure is depicted using a puncture plot, orbit averaging, map-FTLE, and direct topology extraction.

Abstract—Area-preserving maps are found across a wide range of scientific and engineering problems. Their study is made challenging by the significant computational effort typically required for their inspection but more fundamentally by the fractal complexity of salient structures. The visual inspection of these maps reveals a remarkable topological picture consisting of fixed (or periodic) points embedded in so-called island chains, invariant manifolds, and regions of ergodic behavior. This paper is concerned with the effective visualization and precise topological analysis of area-preserving maps with two degrees of freedom from numerical or analytical data. Specifically, a method is presented for the automatic extraction and characterization of fixed points and the computation of their invariant manifolds, also known as separatrices, to yield a complete picture of the structures present within the scale and complexity bounds selected by the user. This general approach offers a significant improvement over the visual representations that are so far available for area-preserving maps. The technique is demonstrated on a numerical simulation of magnetic confinement in a fusion reactor.

Index Terms—Poincaré map, dynamical systems, topology, chaos, area-preserving maps, invariant manifolds.

1 INTRODUCTION

Many natural phenomena can be elegantly described in terms of the Hamiltonian formalism. In areas as different as orbital mechanics, quantum mechanics, and fluid dynamics but also molecular dynamics, chemical dynamics, or ecology the same basic principle of *stationary action* applies. In its simplest form this principle states that the evolution of a system is entirely dictated by the variations of a single scalar function – the *Hamiltonian*. The resulting equations form a dynamical system whose structural properties can be investigated geometrically through *maps*. Maps can be thought of as describing the successive discrete states of an evolving system. Hamiltonian systems in particular give rise to *area-preserving maps*, which exhibit a rich structural picture. From a visual perspective, these maps are quite fascinating since they intertwine fractal topology and regions of chaotic behavior.

Beyond their theoretical appeal, area-preserving maps play a key role in various engineering disciplines. In spatial mission planning or fusion research for instance, the understanding and interpretation of

these maps is required to gain ever deeper insight into the fundamental properties of the considered phenomena. Yet, their intrinsic complexity makes this analysis challenging. While so-called puncture plots offer a straightforward means to obtain a rough picture of the topological features, they are very limited in their ability to offer a reliable picture of the main structures. They yield discrete representations where the main manifolds of the topology can be difficult to identify and some significant patterns can be missed altogether. Despite the introduction in recent years of several techniques to address the shortcomings of these plots, the effective analysis of maps remains a difficult task.

We propose in this paper a new method for the automatic extraction and visualization of the main topological structures present in area preserving maps. Note that in the following we restrict our considerations to maps with two degrees of freedom, though this case covers a wide range of scenarios of practical relevance. In contrast to existing approaches, our algorithm explicitly identifies the location of the fixed points and constructs the manifolds that connect them. Our method offers a precise picture of the island chains, which are the signature of these maps, and thus permits an accurate measure of their quantitative properties, which have an important physical interpretation. Moreover, the reliable extraction of this structural backbone allows us to create compelling visualizations that can be used to quickly identify and interpret the major trends in the data in a transient evolution. To augment the obtained visualizations, we further investigate the visualization of the topological context using *orbit averaging* an adaptation of convolution-based methods such as e.g. *Line Integral Convolution* to the setting of maps.

- Xavier Tricoche is with Purdue University, E-mail: xmt@purdue.edu.
- Christoph Garth is with University of California, Davis, and the University of Kaiserslautern, Germany, E-mail: garth@cs.uni-kl.de.
- Allen Sanderson is with the SCI Institute, University of Utah, E-mail: allen@sci.utah.edu.

Manuscript received 31 March 2011; accepted 1 August 2011; posted online 23 October 2011; mailed on 14 October 2011.

For information on obtaining reprints of this article, please send email to: tvcg@computer.org.




Fig. 2. (a) illustrates two iterations of a Poincaré map. (b) and (c) illustrate islands of resonance in the topology of the Poincaré section. In the integrable case (b), separatrices connect saddle points in the Poincaré map, forming the boundary of an island containing a center point. In the chaotic case (c), the connections are replaced by the intersection of stable and unstable manifolds at an infinite number of points, forming the tangles that characterize chaos. Quasi-periodic orbits exist both inside and outside of the island, where they densely populate KAM manifolds.

We have applied our method in a practical scenario corresponding to a numerical simulation of magnetic confinement in a Tokamak fusion reactor. In this context of *magneto-hydrodynamics* (MHD) simulations, the ability to characterize topological transformations in the early stages of the simulation is key since it provides a crucial insight into the long term evolution of the system (loss of stability, loss of confinement, etc.). It is important to note that we restrict our considerations to *near-integrable systems*. In other words our method is not meant to process fully stochastic systems. Indeed, the basic premise of our method is that the most significant features of the map can be captured through relatively simple geometry, an assumption that is no longer valid if the system is dominated by chaos.

The contents of this paper are organized as follows. Basic definitions and theoretical results relevant to the presentation of our method are first introduced in Section 2. We then briefly review previous work on map visualization and analysis in Section 3 before describing the algorithmic details of our new method (Section 4) and numerical considerations (Section 5). Visualization aspects of our solution are addressed in Section 6, where we introduce the orbit averaging approach, and followed by results in Section 7. Specifically, we consider the important special case of the *Standard Map* (Section 7.1) before commenting on our experience with Tokamak simulation data in Section 7.2. Finally conclusion and future research directions are discussed in Section 8.

2 THEORETICAL AND NUMERICAL FOUNDATIONS

In the following, we provide an overview over the mathematical and numerical concepts that form the basis for the present work. Hamiltonian systems have received considerable theoretical and numerical attention in mathematics, physics, and computational science. Our modest goal in this section is to summarize the key mathematical concepts and results that are needed in the description of the proposed numerical and algorithmic strategies. The interested reader is further referred to excellent classical references [23, 29].

2.1 Hamiltonian Systems and Area-Preserving Maps

Many systems of mechanics (and optics) can be described by a simple set of ordinary differential equations, known as Hamilton's equations:

$$\frac{dq_i}{dt} = \frac{\partial H}{\partial p_i} \quad \frac{dp_i}{dt} = -\frac{\partial H}{\partial q_i}. \quad (1)$$

The *state* of the system at time t is therefore entirely described by a point, $\mathbf{z} = (p_1, \dots, p_N, q_1, \dots, q_N)$, in the $2N$ -dimensional *phase space*. The p_i and q_i are called *momenta* and *positions*, respectively, and the (scalar) function $H(\mathbf{p}, \mathbf{q}, t)$ is called the *Hamiltonian*. In physical systems, the latter typically describes the total energy of the system.

Expressing the evolution of a system in terms of a general ordinary differential equation (ODE) initial value problem:

$$\frac{d\mathbf{z}}{dt} = \mathbf{f}(\mathbf{z}), \quad \mathbf{z}(0) = \mathbf{z}_0 \quad (2)$$

whose solution we denote $\mathbf{z}(t, \mathbf{z}_0)$, we can define its associated *flow map* $\{\phi_t\}_{t \in \mathbb{R}}$

$$\phi_t(\mathbf{z}_0) := \mathbf{z}(t, \mathbf{z}_0), \quad \phi_0(\mathbf{z}_0) \equiv \mathbf{z}_0. \quad (3)$$

The flow map therefore describes the mapping induced by the dynamical system. Equivalent to Equation 1, the *variational principle*, or principle of least action, states that the flow map is constrained to curves, \mathcal{C} in phase space along which the *action integral*:

$$\int_{\mathcal{C}} \mathbf{p} \cdot d\mathbf{q} - H dt$$

is maximized. If the Hamiltonian H itself is an invariant of the motion, the system is said to be conservative.

2.2 Poincaré Map

The *Poincaré map* is a fundamental tool in the study of dynamical systems exhibiting periodicity. Restricting our considerations to a Hamiltonian system with two degrees of freedom (*i.e.* $\mathbf{z} = (p_1, p_2, q_1, q_2)$), a case that covers the application scenarios considered in this work) we first observe that for a given value of the Hamiltonian $H = E$, we can express one of the variables in terms of the others, say $p_2 = p_2(p_1, q_1, q_2, E)$, and study the system in a 3D coordinate system (p_1, q_1, q_2) , where the motion is confined to a doughnut-shaped *invariant torus* also known as *energy surface*. We then construct the Poincaré map by first selecting a Poincaré section Π , in other words a plane that is transverse (*i.e.* nowhere tangent) to the flow, say $\Pi = \{q_2 = 0\}$. A point on the plane is therefore described by its coordinates $\mathbf{x}(x, y) := (p_1, q_1)$. By following the trajectory from this point, we define the Poincaré or return map, \mathcal{P} , via $\tilde{\mathbf{x}} = \mathcal{P}(\mathbf{x})$, where $\tilde{\mathbf{x}}$ is the first intersection of the trajectory emanating from \mathbf{x} and the plane Π , see Figure 2(a). Observe that the fact that the trajectories return to a vicinity of the starting point (*quasi-periodic* behavior) is a consequence of the fact that the energy surface is bounded.

An essential property of Hamiltonian systems compared to other dynamical systems is that the volume of a transported region of the phase space is preserved by the flow map. As a consequence, the Poincaré map itself is *area-preserving* and \mathcal{P} is divergence-free: $\nabla \cdot \mathcal{P} = 0$. It is important to note that the Poincaré map induces a discrete dynamical system in the Poincaré section in which the motion is described in terms of recurrence relations of the form

$$\mathbf{x}_{n+1} = \mathcal{P}(\mathbf{x}_n).$$

In particular, \mathcal{P} is not a vector field.¹

2.3 Integrable, Ergodic and Chaotic Motion

The simplest picture exhibited by Hamiltonian systems correspond to the so-called *integrable* case, in which the motion is completely ordered: the orbits $\mathbf{z}(\cdot, \mathbf{z}_0)$ are either closed (and therefore periodic) or they are confined to tori that are invariant under the flow map. In the Poincaré section these tori appear as nested closed curves.

¹Note for instance the orientation discontinuity at the intersection point of stable and unstable manifolds in chaotic tangles.

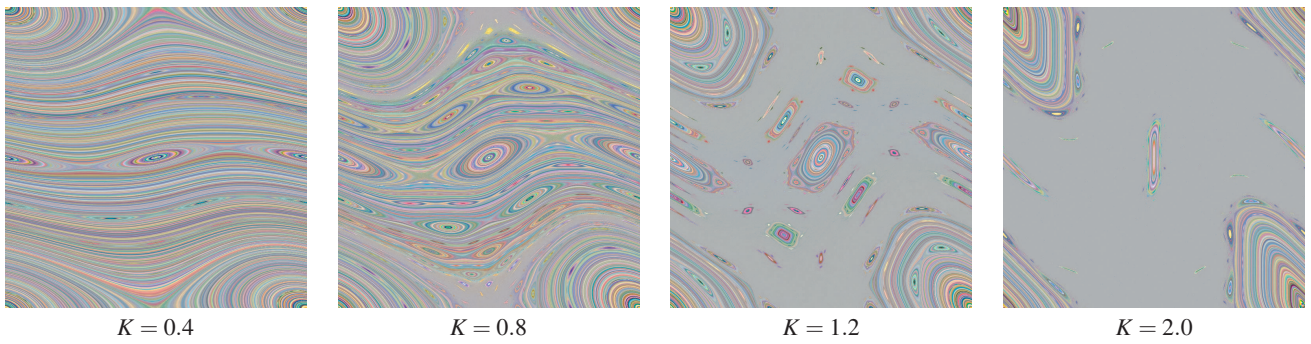


Fig. 3. Evolution of the Standard Map (described in Section 7.1), illustrated using *orbital averaging*. Fixed points and quasi-periodic KAM orbits are identified by identical color averages. Growing ergodic regions appear as uniformly grey areas.

In the opposite limit of *ergodic* behavior the motion is essentially random. In contrast, *chaotic* systems are neither completely integrable nor completely ergodic. Chaotic systems are typical in practice and the primary challenge posed by their analysis is to understand how the structure of phase space progressively “breaks” as the system deviates from integrability. In particular, so-called islands progressively form in the phase portrait in which a locally linear motion develops. Further, irregular trajectories emerge that wander across circumscribed regions of phase space called *ergodic sea*. Their motion is such that they come arbitrarily close to any position within those regions and thus a single trajectory suffices to densely populate the corresponding portion of the Poincaré map. Figures 3 and 4 illustrate this for the *Standard Map*, which we describe in Section 7.1.

2.4 Periodic Orbits and Linearized Motion

Periodic orbits of the system are fixed points of the Poincaré map for a given period. A periodic orbit of period p is invariant under p iterations of the Poincaré map:

$$\mathcal{P}^p(\mathbf{x}_0) := \mathcal{P}(\mathcal{P}^{p-1}(\mathbf{x}_0)) = \mathbf{x}_0,$$

where p is the smallest positive integer that satisfies this relation. We refer to \mathbf{x}_0 as a fixed point of period p . The nature of a fixed point can be determined by a local linear analysis of the Poincaré map in its vicinity. More precisely, a linear approximation of the local motion about \mathbf{x}_0 is based on the Jacobian $\mathbf{J}_p := \nabla_x \mathcal{P}^p$:

$$\mathcal{P}^p(\mathbf{x}_0 + \mathbf{dx}) = \mathcal{P}^p(\mathbf{x}_0) + \mathbf{J}_p(\mathbf{x}_0)\mathbf{dx} + \dots = \mathbf{x}_0 + \mathbf{J}_p(\mathbf{x}_0)\mathbf{dx} + \dots$$

and the eigenvalues λ_i , $i = \{1, 2\}$ of \mathbf{J}_p determine the nature of the fixed point. If they are complex conjugates then the Poincaré map will display elliptical motion under the mapping near the fixed point. This “O-point” configuration is often called *center* in the visualization literature [20]. A local “island” of dominantly regular motion will exist (Figure 2(b)). If the eigenvalues are real and of opposite sign, \mathbf{x}_0 forms a *saddle* (or “X-point”) and the eigenvector of \mathbf{J}_p associated with the negative (resp. positive) eigenvalue aligns with the so-called *stable* (resp. *unstable*) manifolds of \mathbf{x}_0 . These *invariant manifolds* form the boundary of the islands and are often referred to as *separatrices* of the topology. Saddle points (and the associated unstable orbit) lie at the heart of chaos. In general, stable and unstable manifolds intersect in patterns called *chaotic tangles*, and their structure dominates the dynamics in the chaotic sea, see Figure 2(c), bottom left.

Periodic orbits, and in particular their stability, provide crucial information about the local phase space topology. The periodic orbits of most interest are the so-called well-ordered periodic orbits that are continuously connected (in perturbation parameter) to the periodic orbits existing in the integrable limit (i.e. not those periodic orbits which bifurcate out during the chaos).

2.5 Quasi-periodic orbits and KAM theory

In addition to fixed points, islands, and ergodic seas, the Poincaré map reveals the presence of manifolds that are densely covered by quasi-periodic orbits, forming curves on the Poincaré section. The period of

these orbits is measured as their *irrational* winding number, which is defined by following equation.

$$q = \lim_{n \rightarrow \infty} \frac{2\pi n}{\theta(n)}, \quad (4)$$

where n corresponds to the number of revolutions completed around the major axis of the torus and $\theta(n)$ denotes the cumulated angle of the corresponding rotation around the minor axis of the torus. The fundamental *Kolmogoroff-Arnold-Moser* (KAM) theorem [21, 2, 30] states that those manifolds (so-called *KAM surfaces*) that have “sufficiently irrational” winding numbers will survive the onset of chaos through nonlinear perturbations of the integrable configuration.² Here, the degree of irrationality of an irrational number q corresponds to the length of the *continued fractions* [31]

$$q = a_0 + 1/(a_1 + 1/(a_2 + 1/(a_3 + \dots$$

that is required to approximate q with a given ε precision. The KAM surfaces form perfect barriers to transport in the phase portrait, hence their crucial importance in confinement problems. As the degree of chaos increases these surfaces are progressively destroyed. The study of transport barriers in Hamiltonian systems, as visualized by the topology of the Poincaré map, is a crucial aspect of the analysis of magnetic fields in Tokamak fusion reactors since they influence and determine plasma movement. We investigate this application in Section 7.2.

3 PREVIOUS WORK

While discrete dynamical systems and area-preserving maps are not a common topic in visualization publications, a number of previous contributions made in various research fields provide the foundations of our method. We briefly summarize them next.

Topological approaches were introduced in visualization by Helman and Hesselink [19] who first demonstrated the ability of the topological skeleton to offer a schematic and insightful picture of a (continuous) flow. Numerous contributions have since been made to that general methodology and it remains an active research area [17, 18]. In particular, this methodology has been applied to 3D vector fields [41], parameter-dependent vector fields in 2D and 3D [45, 40, 10, 42, 48], and nonlinear singularities [38, 44, 47]. It has also been extended to symmetric tensor fields in 2D [8] and 3D [52], and asymmetric tensor fields [51, 50]. Finally, a closely related topological formalism rooted on Morse theory and combinatorial principles has been investigated for the analysis of complex scalar fields [5, 12, 13, 46, 14].

Most of corresponding papers are concerned with scalar, vector, or tensor *fields*, and are therefore focused on *continuous* systems. A notable exception is the work by Löffelmann et al. who created several methods for the intuitive visualization of discrete dynamical systems

²Note that KAM surfaces, like stable and unstable manifolds of saddle points, constitute invariant manifolds, in other words, low-dimensional manifolds from which the map cannot escape.

defined analytically [26, 24, 25]. In particular, they devised representations that aim at revealing the continuous structures underlying the map. Later, Peikert and Sadlo applied a Poincaré map approach to the visualization of vortex rings in a flow recirculation bubble [33, 34]. While the resulting map is not strictly area-preserving, they showed that the topological structures that are visible in this context share many similarities with them. Following this basic observation, they proposed an image-based method to construct the separatrices that originate from the saddle points located at each extremity of the recirculation bubble; this allowed them to reveal the convoluted patterns formed by their successive intersections as well as the associated island chains and stochastic behavior. Additionally, they described an iterative scheme to compute the *O-points* located in each of the main islands by successive approximation of the location rotation pattern of the map. Most recently, we applied the concept of finite-time Lyapunov exponent [16] to the image-based visualization of invariant manifolds in area-preserving maps [43]. In contrast to the present work however, the proposed technique did not support the geometric extraction or explicit characterization of the map topology. Instead the manifolds were shown as soft edges in the computed picture.

Hamiltonian maps have also been studied in artificial intelligence and data mining. In his Ph.D. thesis, Yip developed a computer system (named *KAM* after the theorem, see Section 2.5) able to navigate autonomously maps defined by analytical expressions [49]. The basic idea behind his approach is to combine geometric and graph theory criteria commonly used in artificial intelligence to permit the automatic recognition of the three main types of orbits that are typically present in a map: closed loop, island chain, and separatrix. Hence, his algorithm explores the map by forming high-level hypotheses about the topology responsible for the observed local patterns. By following a similar approach, Bagherjeiran and Kamath applied a data mining approach to identify patterns in Poincaré plots acquired experimentally [3]. Because of the nature of this type of data, their method assumes no a priori knowledge of the points ordering. Instead, a minimum spanning tree followed by a clustering step is used to infer their 1-dimensional structure. Both of these approaches are essentially suited for the detection and classification of rather large structures of the map. Additionally they do not provide the explicit boundaries of the structures but rather aim at detecting the main features.

In the specific context of fusion reactor simulations, which we consider in Section 7.2, Sanderson et al. described a method to visualize the discrete points of a Poincaré plot by representing them through a




Fig. 4. Close-up in phase portrait of standard map at $K = 1.5$. The fractal nature of the corresponding topology is visible.

curve that connects points lying on the same tori or islands [37, 36]. To do so, the irrational period of the manifold is first approximated based on the available iterates of the map and is used to determine the best connection. The available curve geometry can then be leveraged to approximate the location of the *O-points*, similar to the technique proposed by Peikert and Sadlo [34]. The visualization is built by probing the map at a discrete set of locations and takes advantage of inherent symmetries in the Tokamak.

In addition, researchers in applied mathematics have also considered maps from an algorithmic perspective. In particular, England et al. proposed a method to construct stable and unstable manifolds in Poincaré maps from saddle points by successively extending the piece of manifold that has already been computed. This algorithm forms the basis of the solution we describe in Section 4. Finally, the method describe hereafter makes heavy use of the fundamental notion of Poincaré index [1] to determine the presence of a singularity in a given region. A germane idea was previously used by Mann and Rockwood to identify critical points in nonlinear 3D vector fields [27].

4 TOPOLOGY EXTRACTION

We present in this section a novel algorithm for the extraction of the topological structure of chaotic area-preserving maps that is applicable to both analytically defined Hamiltonian flows and (nearly) divergence free simulation datasets. Emphasis is put on robustness and computational efficiency.

4.1 Constraints

It follows from the discussion in Section 2.3 that the analysis of area preserving maps depends heavily on an accurate and efficient integration of the flow map $(\phi_t)_{t \in \mathbb{R}}$. This computation is required to obtain iterates of the Poincaré map $\mathcal{P}^p, p = 1, \dots, p_{\max}$, where p_{\max} is some upper bound that depends on the goals of the analysis. An exception to this rule are discrete analytical maps where an explicit formula \mathbf{f} describes the recurrence $\mathbf{x}_{p+1} = \mathbf{f}(\mathbf{x}_p)$. Such maps are of academic interest and we consider one such map in Section 7.1. In general, however, the computation of the Poincaré map is made challenging by the need to maintain long term accuracy in the numerical integration of an ODE. In the context of Hamiltonian systems in particular, the property of area-preservation is essentially impossible to guarantee through conventional integration schemes such as Runge-Kutta methods [28]. So-called geometric (or symplectic) integrators do explicitly enforce the invariance of these properties along the integration [15]. However, their application requires a specific formulation of the dynamics (e.g., an explicit expression for the Hamiltonian of the problem), which is not always available. Further, when processing numerical simulations (such as the one considered in Section 7.2), the output data and its continuous reconstruction through piecewise polynomial functions is not exactly conservative. In this case, the area-preserving property is a theoretical reference for the behavior of the studied phenomenon rather than a numerical reality.

Finding fixed points in Poincaré sections in this setting is a computationally challenging and numerically ambiguous task. Map iteration by numerical solution of ODEs at high precision is costly, and there is little hope resolving *all* structures due to the fractal nature of the considered maps. Essentially, looking deep enough, there are structures *everywhere* in the partially chaotic case that we are interested in here. Furthermore, at fixed numerical resolution, an actual fixed point can be impossible to distinguish from a rational surface. Here, as done in previous work, it is typically necessary to introduce parameters in the form of (arbitrary) thresholds that restrict the search spatially strongly influence the analysis.

The basic idea of our approach consists in extracting the fixed points of the map along with the associated island topology sequentially over a prescribed range of periods of the map \mathcal{P} . Restricting the range of periods that should be considered is necessary for computational reasons (the map becomes more and more costly to evaluate as the number of iterations increases) as well as for accuracy concerns. As discussed in the following, the numerics makes it impossible to meaningfully characterize topological structures beyond a certain period

as numerical noise (both in the data and in our processing) starts to dominate the picture. However, this restriction does not penalize the analysis in any significant way since the corresponding topological structures become negligible (in terms of their covering of the phase portrait) as the period exceeds a certain threshold.

For these reasons, the algorithms we contribute in the following emphasize primarily combinatorial principles and consistency checks to address the error and uncertainty that is inherently part of the analysis of such mappings. Our analysis strategy is rooted on key topological principles:

- A non-zero index of the Poincaré map, computed over the boundary of a small region, guarantees that the enclosed region contains at least a fixed point.
- Area-preserving maps exhibit quasi-periodicity, and the KAM manifolds that densely populate the phase portrait in non-chaotic regions have an (irrational) winding number that varies smoothly between neighboring manifolds.
- Fixed points correspond to isolated orbits with rational winding number. Here, winding number q and period p of the map are related by $q = \frac{p}{n}$, where n is the number of rotations along the periodic dimension completed before the map returns to its start position after p iterations.

In the following, we discuss our topology analysis algorithm based on these principles and comment on its properties and parameters.

4.2 Algorithm

Our method proceeds sequentially through six stages that we describe in order of execution. First, we compute a regular sampling of the domain (Step 1), and winding numbers are approximated at every vertex of the corresponding grid (Step 2) that indicate possible periods for topological structures. Subsequently, grid cells that may contain fixed points are identified by evaluating their Poincaré index (Step 3) for the previously determined periods. Candidate cells are then passed to Step 4, where fixed points are located exactly and classified according to their first-order nature into saddles or centers. To ensure a consistent visualization result, in Step 5 we account for fixed points whose presence is required to ensure topological consistency but were been missed by the direct search in Steps 2 through 4. Finally, in Step 6, we compute the separatrices that connect fixed points of saddle type.

The topological structures identified using this algorithm are then visualized as described in Section 6. In the following, we give a detailed description of the individual steps of our method.

Step 1: Initial Sampling

We first perform an initial sampling of the map \mathcal{P}^p for a given range of periods $1 \dots 2p_{\max}$ over the domain of interest using a regular grid of size $N_1 \times N_2$. This bounds the computational effort by limiting analysis to structures that can be reliably resolved at the resulting resolution, and precludes distracting investigation of fractal structures. Practically, the resolution of this initial sampling directly controls the spatial scale of the structures that can be characterized. A coarse sampling will be sufficient to resolve large island chains while a finer sampling will be needed to properly detect skinny islands. In practice, we choose N_1 and N_2 as small multiples (e.g. $\times 1$, $\times 2$, $\times 4$) of the resolution of the computational mesh that represents the system whose Poincaré map we analyze (cf. Section 7.2). This is an intuitive choice for the scientists who are generally wary of analyzing structures beyond the spatial resolution of the computational mesh. In the case of an analytic map representation (such as discussed in Section 7.1), the resolution should be chosen to resolve the smallest features of interest adequately.

The value p_{\max} determines the maximum period of structures to be analyzed. In terms of the scientific application of our method we consider in Section 7.2, higher periods typically correspond to skinnier island structures that are of limited visual and scientific significance,

and analysis in that context is typically focused on lower rational periods. While in principle we could have sampled exactly p_{\max} iterations of the map, we find that doubling this number allows a better assessment of winding number and orbit period, to be determined in the next stage of our algorithm.

Step 2: Period Identification

For every vertex of the discretization, we approximate the (irrational) winding number of \mathcal{P} from the available iterates computed in the previous step. We approximate Equation 4 using the finite sum

$$q(x) = 2\pi \sum_{p=1}^{p_{\max}} \frac{p}{\theta_p} \frac{1}{\|\mathcal{P}^p(x) - x\|}$$

where p designates a number of iterations of the map and θ_p denotes the corresponding rotation angle. In other words, the winding number is obtained as a weighted average of its approximation at each iteration, whereby the weight is chosen inverse proportional to the distance to the seed point. The winding number is generally smooth, and thus this approximation gives good results even for low values of p_{\max} .

To determine possible periods locally, we consider for each cell of the discretization the (irrational) interval $[q_{\min}, q_{\max}]$ spanned by the winding numbers computed at the vertices. Possible map periods p are then identified using the constraint

$$(\gcd(p, p_{\max}) = 1) \text{ AND } (q_{\min} < \frac{p}{p_{\max}} < q_{\max}), \text{ for } p = 1, \dots, p_{\max}.$$

We note all such periods for each cell for further investigation.

Step 3: Poincaré Index Computation

The main strategy we apply to find fixed points is based on the concept of Poincaré index. In planar vector fields, this index measure along a closed curve the number of signed rotations of the field [11]. Isolated linear critical points (where the field is zero) are known to have an index value of ± 1 : saddle points have index -1 , while all other linear critical points have index $+1$. In the case of map, the vector field approximation is valid in the vicinity of a fixed point, by virtue of the linear analysis approach described in Section 2.4. Hence we use the Poincaré index to identify cells that contain a fixed point and compute their exact location using a numerical search.

For each cell C in the discretization of the domain, and for each possible period identified in the previous step, we compute the Poincaré index of the cell by considering the angle integral

$$\frac{1}{2\pi} \int_{\partial C} d\theta(\mathcal{P}^p - Id) \quad (5)$$

along the bounding edges of the cell. We approximate this integral using a piecewise linear representation of the angle along the boundary. However, to accommodate the nonlinearity of \mathcal{P}^p , we adaptively subdivide the edges until the angle between vectors at the vertices of each linear segment falls below an upper bound. To accurately compute the Poincaré index, an upper bound of $\frac{\pi}{2}$ is sufficient and is used in our implementation. If the angle cannot be resolved with a reasonable level of subdivisions, we mark the edge as “ambiguous” and reserve the cell for special treatment later. If Equation 5 yields a non-zero result, the cell is certain to contain a fixed point of period p , and we proceed with locating its exact position in the next phase of the algorithm.

In the case of cells with ambiguous edges, index computation is attempted over an enlarged 3×3 neighborhood of cells. Again, if the index is determined as non-zero, the larger region contains a fixed point and is subject to fixed point location, described next.

Step 4: Fixed Point Location and Linear Characterization

For each of the cells with non-zero index at period p , we employ Newton’s method with linear relaxation [6] to determine the exact location of the contained fixed point. To determine a good starting location within the cell, we sub-sample

$$\|\mathcal{P}^p - Id\| \quad (6)$$

inside the cell and employ the point that minimizes this expression as the starting location. The Jacobian matrix

$$\nabla_{\mathbf{x}} \mathcal{P}^p - I$$

is computed using central differences. Here, we choose the stencil size h as a fraction of the cell edge lengths in the corresponding directions (10% in our implementation). This choice is commensurate with the size of the cell and implicitly assumes that locally linear structures are expected to arise within each cell at the prescribed sampling resolution. The Newton iteration is terminated if the algorithm determines that Equation 6 cannot be further minimized, thereby achieving robustness with respect to small scale noise inherent in the map evaluation.

After the fixed point location x_0 is determined, two cases are distinguished. If the index of the cell was $+1$, a center was found and the linear characterization is complete. If the index was -1 however we compute the map Jacobian at x_0 from which we derive the eigenvectors needed in the separatrix computation.

Step 5: Topological Consistency

If a fixed point x_0 of period p was identified in the previous step, its map iterates

$$\mathcal{P}^{p'}(x_0) \quad \text{for } p' = 1, \dots, p-1$$

are also fixed points of period p of the same type (saddle or center) and form a chain. If not already identified previously, we note these additional fixed points, and, in the case of saddle nature, compute their eigenvectors. This captures the relevant cases in which the computation of the Poincaré index failed.

Step 6: Computation of Separatrices

After all fixed points have been extracted, the separatrices connecting the saddle points of the same chain must be computed. As discussed in Section 2.3 these curves form the boundaries of islands. In the generic chaotic case, these connections are numerically unstable. They correspond to the one-dimensional intersection of the two separatrices that emanate from neighboring saddle points along the chain. Observe that this configuration is not one addressed by the *saddle-connector method* [41] used in flow visualization, through which streamlines connecting a pair of 3D saddle points are extracted as the intersection of their stable and unstable 2D manifold (*i.e.* separating stream surface), respectively. In that configuration the connection is stable, which is in contrast with the present case.

In contrast, these separatrices intersect at an infinite number of discrete locations (*cf.* Figure 2(c)) creating in between oscillations of growing amplitude the closer one moves towards either hyperbolic fixed point [22]. Hence the goal of our separatrix computation is not to precisely account for these oscillations (which is fundamentally impossible) but to provide a schematic representation of that connection that visually conveys the overall geometry of the island and facilitate an assessment of its width. This latter measure is a typical significance criterion in numerical simulations.

Practically, the method we use to compute separatrices is based on the technique proposed by England et al. [9]. The basic idea of their method consists in constructing the separatrix through the solution of successive shooting problems that iteratively advance its end point. The construction is initialized by selecting a point \mathbf{x}_0 in the immediate vicinity of the saddle point along a considered eigenvector and defining $\mathbf{x}_1 = \mathcal{P}^p(\mathbf{x}_0)$, where p is the period of the saddle. This creates an interval $[\mathbf{x}_0, \mathbf{x}_1]$ in which a point \mathbf{x}_2 is determined such that $\mathcal{P}^p(\mathbf{x}_2)$ lies within some prescribed distance of \mathbf{x}_1 . This procedure is repeated iteratively by selected at step i a point \mathbf{x}_{i+1} such that $\mathcal{P}^p(\mathbf{x}_{i+1})$ remains close to \mathbf{x}_i . Note that the parameters of this computation follow the recommendation of [9].

To tackle the challenge posed by the chaotic tangles (not explicitly considered in [9]) we restrict the range of each separatrix to end at 99% of the distance separating two consecutive saddle points along a chain. This computation is carried out both forward along the unstable direction of one saddle and backward along the matching stable direction of the next saddle. By combining these two curves one obtains a

high quality visualization of the saddle-saddle connection (similar to what is shown in Figure 2(c)), even over long distances, while controlling the impact of chaos on the resulting geometry. Note that each p -chain of saddles leads to the computation of $4p$ separatrices.

5 NUMERICAL ASPECTS

Before we describe the extraction algorithm for the topology of area preserving maps, we will briefly discuss numerical aspects of working with such maps.

Poincaré Map Evaluation In our algorithm, we rely on the assumption that the Poincaré map can be computed with high accuracy. While there are ODE solvers that provide an assessment of their global error (in contrast to the *local* error control that typical ODE solvers with adaptive step size apply at each step) we adopted in this work the Runge-Kutta Dormand-Prince DP6(5) method [35] whose dense output provides an excellent balance of accuracy and speed in the computation of the Poincaré map in a computational mesh. Practically, a relative error tolerance of 10^{-8} was found to give good results throughout our experimentation.

Jacobian Approximation A second important computational primitive in our analysis is the measure of the Jacobian $\mathbf{J}_{\mathbf{p}} = \nabla_{\mathbf{x}} \mathcal{P}^p$, as required by the linear analysis of the fixed points of \mathcal{P} and their numerical extraction. A straightforward solution to this problem is to use a central differences approach based on a 4-point stencil

$$\frac{\partial \mathcal{P}^p}{\partial x_i}(\mathbf{x}) \approx \frac{1}{2h} (\mathcal{P}^p(\mathbf{x} + h_i \mathbf{e}_i) - \mathcal{P}^p(\mathbf{x} - h_i \mathbf{e}_i))$$

on the Poincaré section.

However, the determination of appropriate spacings h_i , $i = \{1, 2\}$ can be challenging due the wide range of scales represented in the structure of \mathcal{P}^p . As the considered period increases, \mathcal{P}^p can significantly distort the phase space by virtue of area preservation: a stretch in one direction is automatically compensated by a similar compression in the other and islands become extremely “thin”. This problem is compounded by the fact that while excessively large values of h_i yield de facto inaccurate results, too small values lead the derivative computation to fit numerical noise.

While we require the Jacobian for both the Newton iteration as well as for the determination of the linear nature of fixed points, the former is robust under approximation errors. To make the computation of J_p more robust for the latter case, we introduce the following scheme based on the finite difference approach. We introduce two spatial scales h_{fine} and h_{coarse} and corresponding Jacobians J_p^{fine} and J_p^{coarse} . We proceed to determine the linear nature (saddle or center) of the fixed point on both scales; if the result is the same, we employ J_p^{fine} for further analysis. If the results differ, we compute J_p^{mid} using $h_{\text{mid}} = \frac{1}{2}(h_{\text{coarse}} + h_{\text{fine}})$. Then, we select the finer of the two scales for which the Jacobian approximations agree as the relevant scale for further computation. We have found this “voting” scheme to give much superior results with respect to fixed-scale approximation, particularly in the presents of thin islands and near chaotic regions.

For our purposes, the initial selection of scales is tied directly to the resolution of the discrete sampling of the Poincaré map described above in Section 4. Essentially, the choice of scales reflects an expectation on the size of the region over which a fixed point induces nearly linear motion. Thus, limiting this quantity from below avoids spurious results from sub-scale noise that is otherwise excluded from the analysis. Specifically, we select h_{coarse} and h_{fine} as 5% and 0.5% of the cell diagonal. Similarly, we employ h_{coarse} to obtain Jacobian approximations used in the Newton iteration above.

We wish to point out that other schemes exist in the literature to estimate derivatives using varying approximation scales, such as *e.g.*, *Richardson extrapolation* [6]. However, our experiments with these techniques indicate that they are susceptible to the noise fitting problem by ultimately fitting at too small scales, resulting in frequent misidentification of the linear nature of fixed points in our algorithm.

6 VISUALIZATION

The topological skeleton extracted by our method offers a valuable schematic representation of the subtle structures present in the map. In particular, it offers a continuous and precise contour of the island chains that are key features in the analysis of numerical datasets. Our experimentation with simulation dataset shows that a representation in parameter space can offer a more intuitive means to assess the geometric properties of the islands chains, which can become difficult to track when they are strongly curved and their aspect ratio is heavily distorted in physical space.

While topology provides a crisp and objective image of the structural properties, we have found that the addition of scalar maps provides context to the topological visualization. We essentially consider two variants of scalar context. First, we make use of the adaptation of the *Finite-Time Lyapunov Exponent* [16] to maps that we introduced recently [43]. The basic idea underlying this approach is based on measuring the hyperbolicity of particle trajectories is measured with respect to their intersection with the Poincaré section rather than for a given integration length. A direct visualization of the resulting scalar field offers a very intuitive and expressive context to the stable and unstable manifolds and complements them visually. Here, we use the maximum of both forward-time and backward-time FTLE, as obtained from the iterates of \mathcal{P} , to color-code both stable and unstable manifolds in a single scalar field.

As a second scalar field, we consider a technique similar to *Line Integral Convolution* [7] applied in vector field visualization. By overlaying the domain of interest with a noise texture $noise(x)$, we compute

$$a(x) := \frac{1}{p_{\max}} \sum_{p=0}^{p_{\max}} noise(\mathcal{P}^p(x)) \quad (7)$$

for every pixel x in an image, thus essentially averaging the noise texture over a number of iterations of the Poincaré map. This *orbit averaging* procedure results in an image that identifies pixels on the same orbit with identical values of a . In essence, we replace the convolution along streamlines used in LIC by convolution along orbits using a box filter. While it is immediately apparent that this approach assigns the same color to all fixed points on a chain, quasi-periodic orbits become visually apparent by virtue of their dense nature: if p_{\max} is chosen sufficiently large, all pixels on a quasi-periodic orbit are assigned similar values of a . Furthermore, if $noise$ is chosen as uniformly distributed, the value of a converges to chaotic orbits converges to the expectation value of a , e.g. to 0.5 if $noise$ is uniform over $[0, 1]$.

In practice, choosing large values of p_{\max} can be computationally prohibitive due to the large number and overhead of orbit computation. However, several factors enable an efficient computation of this approach. First, as already described in Section 4, a single orbit touches many pixels of a , and thus the computation can be accelerated by reusing the convolution result over a fixed range of iterations in similar manner as *FastLIC* [4]. Second, we can reduce the number of map evaluations by storing the orbit generated by every pixel for a smaller maximal period P_{\max} , and then convolving the image multiple times, each time reusing the output of the previous pass as the input of the current pass. Naturally, high-pass filtering is employed after every pass to preserve image contrast [32]. In practice, we find that by choosing p_{\max} as a small multiple ($2\times$ or $4\times$) of the highest period of interest and then performing several convolution passes results in smooth images that reliably identify interesting structures.

Figure 3 illustrates our approach on the Standard Map (cf. Section 7.1) using uniformly distributed color noise to provide an overview of the map. Here, the averaging results in well-defined (quasi-)periodic structures; chaotic regions appear uniformly grey and are thus easily distinguished. Figures 6 and 8 show our method applied to a Tokamak simulation (cf. Section refsec:tokamak), using uniformly gray noise to provide an illustration of the KAM curves as context to the (primary) topological visualization.




Fig. 5. A simplified schematic of a Tokamak fusion reactor showing its basic components.

7 RESULTS

7.1 Standard Map

The Standard Map (also known as *Chirikov-Taylor Map*) is an area-preserving 2D map of the 2π square onto itself defined as follows:

$$p_{n+1} = p_n + K \sin(\theta_n) \quad (8)$$

$$\theta_{n+1} = \theta_n + p_{n+1}, \quad (9)$$

where p_n and θ_n are taken modulo 2π . K is a parameter that controls the nonlinearity of the map. The Standard Map describes the dynamics of several mechanical systems (e.g., a kicked rotor) and has attracted the attention of theoretical and computational research alike since it is a simple yet powerful tool to study Hamiltonian chaos. In the context of this work, the Standard Map offers a compelling means to test our proposed method across a range of configurations. Clearly, the fixed points of this map can be obtained analytically. Straightforward calculus shows for instance that the fixed points under a single iteration of the map are given by $(p, \theta) = (0, 0)[\pi]$. However our goal here is to investigate the structural properties of this map from a purely algorithmic standpoint, as a proof of concept. In particular, all computations (including the map Jacobian) are performed numerically without resorting to the available analytic expressions. As such this offers a truly interesting test case where a variety of topological configurations are encountered.

Figure 7(a) illustrates the topology of the standard map as extracted by our method in the period range 1 to 20. A close-up view, provided in Figure 7(b), reveals how subtle structure are properly captured by our method despite their challenging distortion. In both figures, context is provided by the FTLE scalar field described in Section 6. Generally, our algorithm is able to extract all structures in the considered period range.

7.2 MHD Simulation of Plasma Confinement in a Tokamak Fusion Reactor

Magnetic fusion reactors, such as the International Thermonuclear Experimental Reactor (ITER), a Tokamak reactor which is scheduled for completion in 2018, are being investigated intensively as a source for future low cost power. In their basic operation, magnetic confinement fusion uses the electrical conduction of the burning plasma to contain it within magnetic fields, see Figure 5. The plasma containment is essential for the efficient functioning of such reactors, since contact of the plasma with the walls of the Tokamak is strongly detrimental to its material integrity.

A critical characteristic of a typical fusion reactor is the growth of instabilities in the plasma due to the large gradients of density and temperature, the field geometry, and the inherent self-consistent interactions between charged particles and electromagnetic waves. Plasma instabilities occur on very different spatial and temporal scales and can represent highly unique phenomena. One such instability, magnetic reconnection, prevents the magnetic field from confining the plasma and

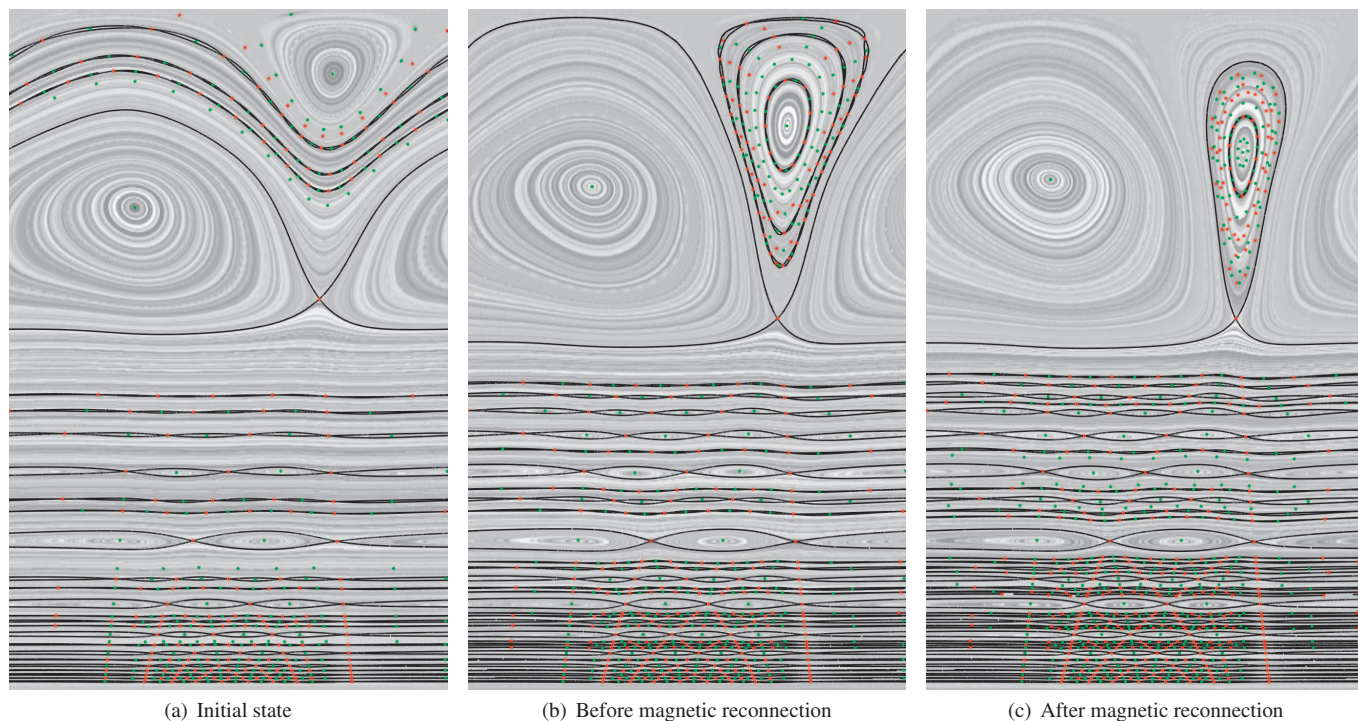


Fig. 6. Topological structure of the Poincaré map of the MHD Tokamak simulation visualized in parameter space. The islands that were detected in the course of the analysis are depicted directly, and context is provided by KAM structures made visible using orbit averaging. The dramatic topological changes during the magnetic reconnection are readily apparent from our combination of techniques. Specifically, the phenomenon of magnetic reconnection can be seen taking place progressively, leading to a configuration where the separatrices of the period-1 saddle point now enclose the period-1 center point. Color coding: red: saddles, green: centers

leads to its transport. Locating and understanding this phenomena can best be done through visualizing the topology of the magnetic field and identifying features within it. The data we study here is a simulation of magnetic reconnection performed using the NIMROD [39] code. The magnetic field is represented on a structured mesh of size $121 \times 81 \times 51$ and available on a small number of time steps extracted from a longer simulation run. Along the first two axes, representing poloidal slices of the Tokamak, bi-quadratic Lagrangian interpolation is employed, while the third axis represents coefficients in a Fourier series that approximate the field along the periodic toroidal direction. As such, reconstructing the vector field and integrating magnetic field lines with high precision is laborious.

In the normal operation of a tokamak reactor, the magnetic field lines are topologically distinct from each other and form a series of concentric flux surfaces that confine the plasma. Because the magnetic field lines are periodic, the topology can clearly be seen by creating a puncture plot, see Figure 1, left. Here, the Poincaré section Π coincides with a poloidal cross section of the torus-shaped tokamak. In the presence of instabilities, the magnetic field can become distorted and form magnetic islands, and it is this formation of islands that is magnetic reconnection. Surrounding these islands are two separatrices that divide the magnetic field in to four domains. At the intersection of these two separatrices lies a single saddle point (aka a separator) that represents a magnetic field line that is truly periodic. It is at this location where plasma transport can occur.

Locating these features; islands, separatrices, and X points, is an important component in understanding plasma transport in magnetic fusion research. However, generating Poincaré plots with sufficient detail is computationally expensive; moreover, unless seed points for the plot are selected in a robust fashion, features within the magnetic field may be missed. As such developing a robust technique that allows for the rapid visualization and facilitates the analysis of topological structures of magnetic field lines in an automatic fashion will aid in the future design and control of Tokamak reactors. In Figure 1,

right, we show the application of our method to time step of a transient numerical simulation of magnetic confinement on a poloidal slice. In contrast, for the sake of visual clarity, and as is common in this context, we consider hereafter the Poincaré map in an (ϕ, r) parameter space mapping to the poloidal slice. Thus, the resulting image is periodic on the horizontal axis which corresponds to the angle in the poloidal plane.

The transient evolution of the topological structures is shown in Figure 6, in which different time steps of the simulation are shown.

One of the more interesting phenomena within magnetic confined fusion is the formation of “islands within islands”. During the latter stages of island growth during the onset of magnetic reconnection two island chains may grow sufficiently large that they overlap and interact with each other leading to the surfaces breaking up into smaller islands. At this point in time the magnetic field is entering into the non-linear stage of magnetic reconnection, and this represents one of the last steps before the magnetic field becomes chaotic. The breakdown of the surface into islands within islands is believed to further increase plasma transport. Figure 8 illustrates a single island within a chain of islands that has broken into a series islands-within-islands. The ability to detect such phenomena using our technique demonstrates its ability to resolve fine scale features that might otherwise be missed easily.

7.3 Performance

The performance of our algorithm is difficult to quantify, since the runtime depends to a significant extent on the quantity and quality of structures identified. All stages of our implementation leverage the embarrassingly parallel nature of the computation (initial regular map sampling, cell index computation, per-cell fixed point extraction, per-chain separatrices construction). We have therefore implemented our method using openmp and observed a linear speedup with respect to the number of available cores, up to 32 cores (8 Xeon Quadcore processors). The entire processing of any given Tokamak time step took under 5 minutes on 32 cores. Note that a mapping of this computation

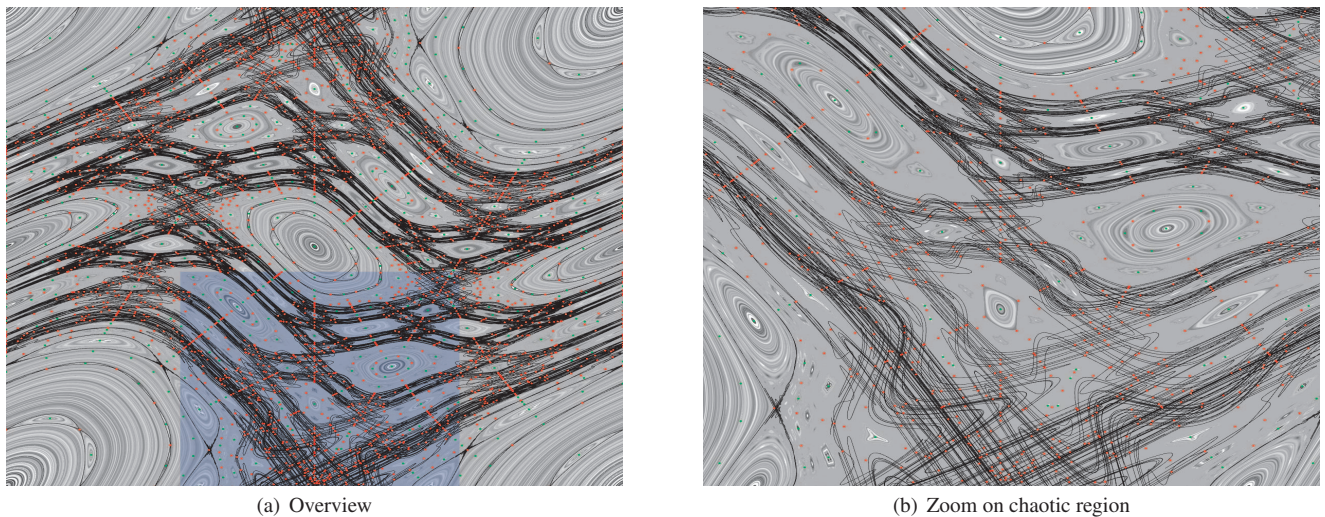


Fig. 7. Topology of the standard map for $K=1.1$. At that stage the map exhibits sizable chaotic regions. However our algorithm is able to capture the most significant structures in the individual islands of the map. A very large numbers of fixed points are identified within the forming chaotic sea, which cannot be reliably linked to separatrices since the amplitude of chaotic motion is maximized along those separatrices. While these points appear random, note the many degrees of symmetry which they exhibit.

to the GPU, while conceptually straightforward and promising from a performance standpoint remains problematic due to possible inaccuracies in the computation (single float precision, limited precision of texture coordinates).

8 CONCLUSION

We have presented an algorithmic and computational framework to permit the automatic topological analysis of area-preserving maps associated with Hamiltonian systems. While these maps are of great theoretical interest they are also very important in practice since they offer a geometric interpretation of the qualitative behavior of complex physical systems. Our approach significantly improves on previous work by allowing for the explicit geometric characterization of very subtle structures that would typically be missed through Poincaré plot investigation of the map. Our algorithm was carefully designed to be numerically robust in a context where chaos and fractal complexity make any numerical processing fundamentally challenging. By restricting our computation to a small number of iterations of the period from any given point (commensurate with the period range relevant to the analysis) we are able to obtain reliable results that are further enhanced by various correction mechanisms driven by topological considerations. From a visualization perspective we have proposed to combine this schematic topological information with the dense and effective visual representation afforded by the concept of orbit averaging, which we introduced in this paper.

We have tested our methods on a standard analytical map and on a transient numerical simulation of magnetic confinement. Our results underscore the potential of our method to effectively support the offline analysis of large simulation datasets, a context in which they can offer a valuable diagnostic tool. In that regard there are many promising avenues for future work. In particular, the transient nature of the considered phenomena is so far handled in a discrete manner, which gives only indirect insight into the topological transformations that control the development of the structures observed in individual plots. In addition the ergodic behavior of field lines close to the boundary of interesting structures constitutes a significant challenge that further research should investigate to ensure well defined island boundaries.

ACKNOWLEDGMENTS

The authors wish to thank S. Kruger, E. Held, C. Sovinec, and J. Breslau for sharing their expertise on MHD and plasma confinement. The numerical data considered in the paper was produced by the software NIMROD (<https://nimrodteam.org/>). The authors are

also indebted to Stuart Hudson for providing his insight in nonlinear dynamics. This work was supported in part by the Director, Office of Advanced Scientific Computing Research, Office of Science, of the U.S. Department of Energy under Contract No. DE-FC02-06ER25780 through the Scientific Discovery through Advanced Computing (SciDAC) programs Visualization and Analytics Center for Enabling Technologies (VACET), the DOE SciDAC Fusion Scientific Application Partnership, the National Science Foundation under grant IIS-0916289, and a gift by Intel Corporation.

REFERENCES

- [1] A. A. Andronov. *Qualitative Theory of Second-Order Dynamic Systems*. John Wiley & Sons, 1973.
- [2] V. I. Arnold. Proof of A. N. Kolmogorov's theorem on the preservation of quasiperiodic motions under small perturbations of the Hamiltonian. *Russ. Math. Surv.*, 18(5):9, 1963.
- [3] A. Bagherjeiran and C. Kamath. Graph-based methods for orbit classification. In *Proc. of Sixth SIAM International Conference on Data Mining*, April 2006.




Fig. 8. Closeup of a single island within a chain of islands that has broken into a series of islands-within-islands and is part of the nonlinear stage of magnetic reconnection.

- [4] H. Battke, D. Stalling, and H.-C. Hege. Fast line integral convolution for arbitrary surfaces in 3d. In S. Berlin, editor, *Visualization and Mathematics*, pages 181–195, 1997.
- [5] P.-T. Bremer, H. Edelsbrunner, B. Hamann, and V. Pascucci. A topological hierarchy for functions on triangulated surfaces. *IEEE Trans. Vis. Comput. Graph.*, 10(4):385–396, 2004.
- [6] R. L. Burden and J. D. Faires. *Numerical Analysis*. PWS-Kent, Boston, 1993.
- [7] B. Cabral and L. Leedom. Imaging vector fields using line integral convolution. *Computer Graphics (SIGGRAPH '93 Proceedings)*, 27(4):263–272, 1993.
- [8] T. Delmarcelle and L. Hesselink. The topology of symmetric, second-order tensor fields. In *Proceedings of IEEE Visualization 1994*, pages 140–147. IEEE Computer Society Press, 1994.
- [9] J. England, B. Krauskopf, and H. Osinga. Computing one-dimensional global manifolds of poincaré maps by continuation. *SIAM Journal of Applied Dynamical Systems*, 4(4):1008–1041, 2005.
- [10] C. Garth, X. Tricoche, and G. Scheuermann. Tracking of vector field singularities in unstructured 3D time-dependent data sets. In *Proceedings of IEEE Visualization 2004*, pages 275–280, 2004.
- [11] J. Guckenheimer and P. Holmes. *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. Springer-Verlag, 1983.
- [12] A. Gyulassy, V. Natarajan, V. Pascucci, P.-T. Bremer, and B. Hamann. Topology-based simplification for feature extraction from 3d scalar fields. In *Proc. IEEE Visualization 2005*, pages 275–280, 2005.
- [13] A. Gyulassy, V. Natarajan, V. Pascucci, P.-T. Bremer, and B. Hamann. A topological approach to simplification of three-dimensional scalar functions. *IEEE Trans. Vis. Comput. Graph.*, 12(4):474–484, 2006.
- [14] A. Gyulassy, L. G. Nonato, P.-T. Bremer, C. T. Silva, and V. Pascucci. Robust topology-based multiscale analysis of scientific data. *Computing in Science and Engineering*, 11(5):88–95, 2009.
- [15] E. Hairer, C. Lubich, and G. Wanner. *Geometric Numerical Integration*. Springer-Verlag, 2006.
- [16] G. Haller. Distinguished material surfaces and coherent structures in three-dimensional flows. *Physica D*, 149:248–277, 2001.
- [17] H. Hauser, H. Hagen, and H. Theisel. *Topology Based Methods in Visualization*. Springer, 2007.
- [18] H.-C. Hege, K. Polthier, and G. Scheuermann. *Topology Based Methods in Visualization II*. Springer, 2009.
- [19] J. Helman and L. Hesselink. Visualizing vector field topology in fluid flows. *IEEE Computer Graphics and Applications*, 11(3):36–46, 1991.
- [20] J. L. Helman and L. Hesselink. Visualizing vector field topology in fluid flows. *IEEE Computer Graphics and Applications*, 11(3):36–46, May 1991.
- [21] A. N. Kolmogorov. On the conservation of conditionally periodic motions under small perturbation of the Hamiltonian. *Dokl. Akad. Nauk. SSR*, 98:469, 1954.
- [22] A. Lichtenberg and M. Lieberman. *Regular and Stochastic Motion*. Applied Mathematical Sciences. Springer Verlag, 1983.
- [23] A. J. Lichtenberg and M. A. Lieberman. *Regular and Chaotic Dynamics, 2nd ed.* Springer-Verlag, New York, 1992.
- [24] H. Löffelmann, H. Doleisch, and E. Gröller. Visualizing dynamical systems near critical points. In *14th Spring Conference on Computer Graphics*, pages 175–184, 1998.
- [25] H. Löffelmann and E. Gröller. Enhancing the visualization of characteristic structures in dynamical systems. In *Visualization in Scientific Computing '98*, pages 59–68, 1998.
- [26] H. Löffelmann, T. Kucera, and E. Gröller. Visualizing poincaré maps together with the underlying flow. In H.-C. Hege and K. Polthier, editors, *Mathematical Visualization: Proceedings of the International Workshop on Visualization and Mathematics '97*, pages 315–328. Springer, 1997.
- [27] S. Mann and A. Rockwood. Computing singularities of 3d vector fields with geometric algebra. In *IEEE Visualization Proceedings '03*, 2003.
- [28] J. Marsden and M. West. *Discrete Mechanics and Variational Integrators*. Acta Numerica, 2001.
- [29] J. E. Marsden and T. Ratiu. *Introduction to Mechanics and Symmetry*. Texts in Applied Mathematics vol. 17. Springer-Verlag, 2003.
- [30] J. Moser. On invariant curves of area-preserving mappings of an annulus. *Nachr. Akad. Wiss. Göttingen, Math. Phys. Kl. II 1,1, Kl(1):1*, 1962.
- [31] I. Niven. *Irrational Numbers*. The mathematical association of America, 1956.
- [32] A. Okada and D. L. Kao. Enhanced line integral convolution with flow feature detection. In *Proceedings of IS&T/SPIE Electronic Imaging*, 1997.
- [33] R. Peikert and F. Sadlo. Visualization Methods for Vortex Rings and Vortex Breakdown Bubbles. In A. Y. K. Museth, T. Möller, editor, *Proceedings of the 9th Eurographics/IEEE VGTC Symposium on Visualization (EuroVis'07)*, pages 211–218, May 2007.
- [34] R. Peikert and F. Sadlo. Flow Topology Beyond Skeletons: Visualization of Features in Recirculating Flow. In H.-C. Hege, K. Polthier, and G. Scheuermann, editors, *Topology-Based Methods in Visualization II*, pages 145–160. Springer, 2008.
- [35] P. J. Prince and J. R. Dormand. High order embedded runge-kutta formulae. *Journal of Computational and Applied Mathematics*, 7(1), 1981.
- [36] A. Sanderson, G. Chen, X. Tricoche, D. Pugmire, S. Kruger, and J. Breslau. Analysis of recurrent patterns in toroidal magnetic fields. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1431–1440, 2010.
- [37] A. Sanderson, X. Tricoche, C. Garth, S. Kruger, C. Sovinec, E. Held, and J. Breslau. Poster: A geometric approach to visualizing patterns in the poincaré plot of a magnetic field. In *Proc. of IEEE Visualization 06 Conference*, 2006.
- [38] G. Scheuermann, H. Krüger, M. Menzel, and A. P. Rockwood. Visualizing nonlinear vector field topology. *IEEE Trans. Vis. Comput. Graph.*, 4(2):109–116, 1998.
- [39] C. R. Sovinec, T. A. Gianakon, E. D. Held, S. E. Kruger, and D. D. Schnack. Nimrod: A computational laboratory for studying nonlinear fusion magnetohydrodynamics. *Phys. Plasmas*, 10(5):1727, 2003.
- [40] H. Theisel and H.-P. Seidel. Feature flow fields. In *Proceedings of Joint Eurographics - IEEE TCVG Symposium on Visualization (VisSym '03)*, pages 141 – 148. ACM, 2003.
- [41] H. Theisel, T. Weinkauff, H.-C. Hege, and H.-P. Seidel. Saddle connectors - an approach to visualizing the topological skeleton of complex 3d vector fields. In *IEEE Visualization '03*, 2003.
- [42] H. Theisel, T. Weinkauff, H.-C. Hege, and H.-P. Seidel. Stream line and path line oriented topology. In *Proceedings of IEEE Visualization '04 Conference*, pages 321–328, 2004.
- [43] X. Tricoche, C. Garth, and A. Sanderson. Visualizing invariant manifolds in area-preserving maps. In *Proceedings of the Topology-Based Methods in Visualization 2011*, 2011 (to appear).
- [44] X. Tricoche, G. Scheuermann, and H. Hagen. A topology simplification method for 2D vector fields. In *Proceedings of IEEE Visualization 2000*, pages 359–366, 2000.
- [45] X. Tricoche, T. Wischgoll, G. Scheuermann, and H. Hagen. Topology tracking for the visualization of time-dependent two-dimensional flows. *Computer and Graphics*, 26:249–257, 2002.
- [46] G. Weber, P.-T. Bremer, and V. Pascucci. Topological landscapes: A terrain metaphor for scientific data. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1416–1423, 2007.
- [47] T. Weinkauff, H. Theisel, H.-C. Hege, and H.-P. Seidel. Boundary switch connectors for topological visualization of complex 3d vector fields. In O. Deussen, C. Hansen, D. A. Keim, and D. Saupe, editors, *VisSym 2004: Joint Eurographics/IEEE Symposium on Visualization*, pages 183–192, Konstanz, Germany, 2004. Eurographics.
- [48] T. Weinkauff, H. Theisel, H.-C. Hege, and H.-P. Seidel. Topological structures in two-parameter-dependent 2D vector fields. *Computer Graphics Forum*, 25(3):607–616, September 2006. Eurographics 2006, Vienna, Austria, September 04 - 08.
- [49] K. M.-K. Yip. *KAM: A System for Intelligently Guiding Numerical Experimentation by Computer*. MIT Press, 1992.
- [50] E. Zhang, H. Yeh, Z. Lin, and R. Laramée. Asymmetric tensor analysis for flow visualization. *IEEE Transactions on Visualization and Computer Graphics*, 15(1):106–122, 2009.
- [51] X. Zheng and A. Pang. 2d asymmetric tensor analysis. In *Proc. of IEEE Visualization 2005*, pages 3–10. IEEE Computer Society Press, Oct. 2005.
- [52] X. Zheng, B. Parlett, and A. Pang. Topological lines in 3D tensor fields and discriminant hessian factorization. *IEEE Transactions on Visualization and Computer Graphics*, 11(4):395–407, 2005.