Visual Exploration of Circulation Rolls in Convective Heat Flows

A. Frasson[1] M. Ender[1] S. Weiss[1] M. Kanzler[1] A. Pandey[2] J. Schumacher[2]
R. Westermann[1]
[1]Computer Graphics & Visualization Group, Technische Universität München, Germany*
[2]Fluid Mechanics Group, Technische Universität Ilmenau, Germany†

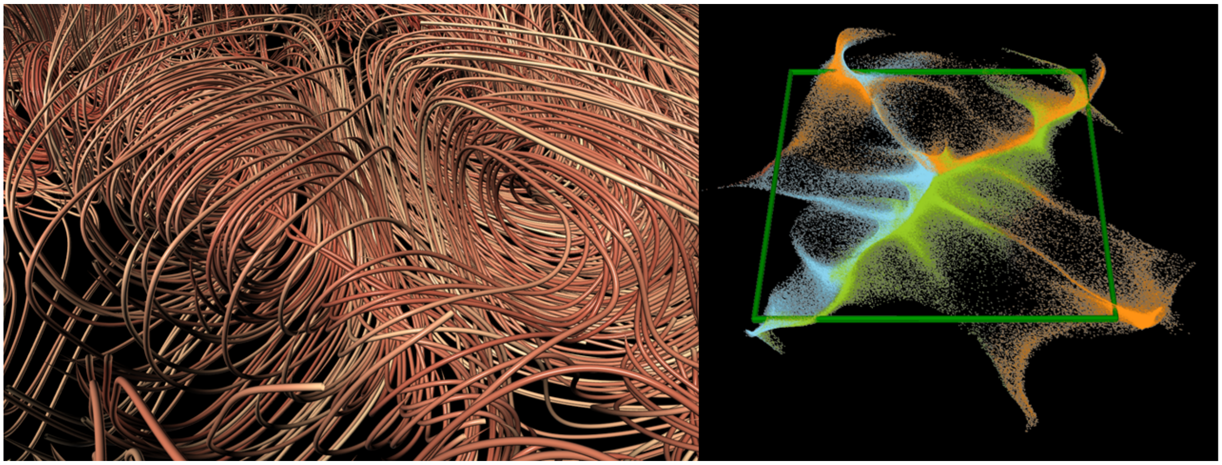

Figure 1: Based on the segmentation of circulation rolls in 3D convective flows, we introduce an importance measure for streamlines to emphasize interior heat flow patterns, and particle-based visualization to analyse the structure and speed of heat mixing across rolls.

Abstract

We present techniques to improve the understanding of pattern forming processes in Rayleigh-Bénard-type convective heat transport, through visually guided exploration of convection features in time-averaged turbulent flows. To enable the exploration of roll-like heat transfer pathways and pattern-forming anomalies, we combine feature extraction with interactive visualization of particle trajectories. To robustly determine boundaries between circulation rolls, we propose ridge extraction in a z-averaged temperature field, and in the extracted ridge network we automatically classify topological point defects hinting at pattern forming instabilities. An importance measure based on the circular movement of particles is employed to automatically control the density of 3D trajectories and, thus, enable insights into the heat flow in the interior of rolls. A quantitative analysis of the heat transport within and across cell boundaries, as well as investigations of pattern instabilities in the vicinity of defects, is supported by interactive particle visualization including instant computations of particle density maps. We demonstrate the use of the proposed techniques to explore direct numerical simulations of the 3D Boussinesq equations of convection, giving novel insights into Rayleigh-Bénard-type convective heat transport.

Keywords: flow visualization, particle-based visualzation, convective heat transport

1 Introduction

Many turbulent flow processes in nature and technology belong to the class of thermal convection flows and are driven by temperature differences. Rayleigh-Bénard convection (RBC) is at the core of all of these examples. In a RBC flow a horizontally extended layer of fluid between two rigid walls is uniformly heated from below and cooled from above. When the temperature difference $\Delta T = T_{bottom} - T_{top}$ across the layer is small, heat is transported only by diffusion from the bottom to the top and the fluid in between is at rest. Once the temperature difference exceeds a critical value, ΔT_c, fluid motion in form of aligned and pairwise counter-rotating rolls sets in. A further increase of the temperature difference triggers bifurcations to more complex roll patterns and eventually a transition to a fully turbulent flow between both plates [4].

At the core of research in turbulent convection is a better understanding of the complex three-dimensional (3D) mechanisms of turbulent heat transport. Although the fluid is fully turbulent, it is not stochastic and featureless. The turbulent transport of heat is organized in form of characteristic coherent structures of the temperature and velocity fields. They are obtained once the turbulent fluctuations are removed by an averaging over a longer time interval with a characteristic time depending on the Rayleigh and the Prandtl number [27]. Fig. 2 shows the effect of time-averaging on the turbulent RBC flow.

Particularly in horizontally extended convection flows, regular and well-ordered patterns have been revealed recently:

- *Thermal plumes*: Blobs of hotter (colder) fluid which rise from the bottom (fall from the top) into the bulk of the fluid layer where they are exerted to turbulent mixing by the existing vortices and disperse.
- *Circulation rolls*: Well separated transport structures in the heat flow which fill the whole convection layer as time-averaged patterns of the velocity field (Fig. 2 (right)).

The plumes form a cellular structure indicating the cells of upward and downward heat transport. The shape and size of thermal plumes has been studied in convection domains of smaller size such as in laboratory experiments with thermochromic particles [50]; their importance for the variability of turbulent heat transport was analysed

*e-mail: firstname.lastname@tum.de
†e-mail: firstname.lastname@tu-ilmenau.de

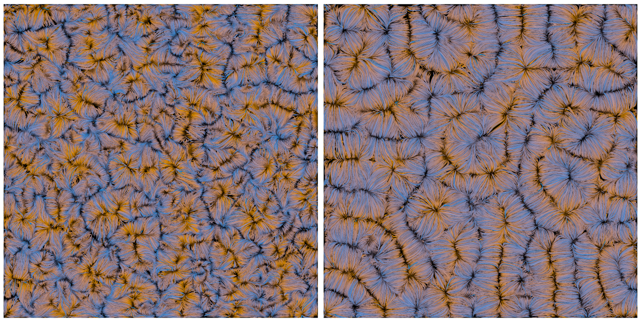

Figure 2: Randomly seeded streamlines in the instantaneous (left) and time-averaged (right) thermal convection flow field. $Ra = 10^5$, $Pr = 0.7$, $\Gamma = 25$. Coloring indicates temperature from low (blue) to high (orange).

in numerical simulations in the Eulerian [39] and the Lagrangian frame of reference [37].

While coherent structures of the temperature field are rather well explored, much less is known on the coherent structures of the 3D velocity field and their interplay. Circulation rolls are usually seen as atomistic heat flow structures with similar statistical properties, and an in-depth analysis of their interior flow structure has not been carried out so far. Circulation rolls are also not independent of each other, and it remains a challenge to analyse the 3D heat transport where rolls interact and heat transport across rolls occur. In particular, it is known that convection flow patterns are never perfectly arranged. Defects in the form of local mergers of two rolls, so-called dislocations, are observed already right above the onset of thermal convection when the system is not turbulent [5].

1.1 Contribution

We introduce a novel combination of 2D imaging and 3D visualization techniques to explore the essential circulation network of the extended convection flow that is hidden by the turbulent fluctuations of the velocity fields. We propose techniques to segment the network and its defects automatically, and use these structures in a subsequent 3D visualization to study the flow across circulation roll boundaries and in the vicinity of defects. The network is used to determine seed locations for streamlines, and to classify lines in the cores of circulation rolls. Streamline visualization using density control in screen-space is then used to uncover the specific heat transport patterns of the roll ensemble. Extracted structures are also used to guide the user to important regions where particles can be seeded. To reveal transport barriers and indicate crossing and mixing events, we use interactive particle tracing and introduce instant computation and visualization of particle density maps. Since both techniques serve different purposes, i.e., visualization of structure via streamlines and of dynamics / statistics via particles, they are complementary to each other and not intended to be used simultaneously. Our specific contributions are:

- A segmentation approach for circulation rolls and defects based on ridge extraction in the z-averaged temperature field.
- An importance criterion for streamlines based on their closeness to the rolls' centerlines, to reveal the fluid flow in the core of circulation rolls.
- A particle-based visualization approach to perform a quantitative analysis of the heat mixing behaviour within and across circulation rolls.

Even though the network segmentation is specifically tailored to the up-down symmetry of RBC flows, our visualization techniques can be applied to any other type of flow where circulation patterns arise and have been segmented. For example, further candidates are the turbulent Taylor-Couette flow between two rotating concentric cylinders [3], or convection-based cell structures caused by atmospheric circulation.

A full application study involving a thorough discussion of the domain-specific interpretation of our results is out of the scope of this paper. To a certain extent, this is because at this point the domain experts were mostly asking for dedicated visualization options to enable a more effective exploration of the complex 3D mechanisms of turbulent heat transport in RBC flows. Since such options, tailored to the specific structure of RBC flows, are not available, in this work we introduce the basic functionality to achieve this.

In the remainder of this paper we first reference previous work that is related to ours and give an overview of the data we use. We continue with a description of how to segment circulation rolls and extract defects using ridges and valleys in the temperature field. Next, we introduce our specific adaptations and extensions to density-controlled streamline visualization and 3D particle-based visualization. We conclude the paper with a discussion of our current work and an outlook on the future applications of the proposed techniques.

2 Related Work

Vector field topology: Our work is related to vector field topology, a powerful tool for analysing the overall structure and dynamics of flow fields, yet it aims specifically at the analysis of advective heat transport in Rayleigh-Bénard type convection. In the visualization community, extensive research has been pursued over the last years on the definition, interpretation, and extraction of topological features, most of them focussing on the analysis of quantities that are solely driven by advection. This includes the exploration of new feature types in stationary 2D and 3D fields as well as time-dependent transport processes [41, 42, 46]. The generalisation of 2D vector field topology to time-dependent and 3D vector fields has been proposed by Sadlo and Weiskopf [32] and Uffinger et al. [44]. Sadlo et al. [30] generalise topological flow analysis to also include diffusion effects, by extracting Langrangian coherent structures in the advection-diffusion fields. Karch et al. [16] stress the importance of numerical diffusion in numerical advection-diffusion simulation and introduce the concept of passive diffusion to visualize diffusion fluxes. Similar to stream tubes. Hochstetter and co-workers [13] provide a new visualization approach, so called stream feathers, to convey advective-diffusive concentration transport.

Much of recent research has been devoted to the concept of the Finite-Time Lyapunov Exponent [11], based upon which a number of extensions regarding the efficient and robust computation have been performed. Especially the extraction of integration-based features as height ridges [7,29,31] is closely related to our work. Specific applications have been addressed, for instance, vortex breakdown bubbles by Peikert and Sadlo [28], magnetic fields [33], and ridge-based extraction of features in the air flow induced by revolving doors [35]. Recently, Machado demonstrated the extraction of bifurcation lines in the space-time representation of a vector field [23].

Line rendering: Our approach for streamline visualization falls into the broader category of techniques that try to find an as small as possible set of lines to represent the important structures in a given line set. Object-space approaches try to adapt the seeding of lines so that uniform line density [25, 40, 43], equally line distribution [14, 20, 21], or coverage of specific flow or line features [26, 45, 48] is achieved. Xu et el. [47] use the local directional variation of a given vector field as entropy measure, and generate an entropy field that is then used to guide the placement of streamlines. Yu et al. [49] use hierarchical clustering to generate a small representative set of streamlines, and select streamlines close to the cluster boundaries to convey the spatial extent of computed clusters.

Screen-space approaches determine for each new view the subset of lines to be rendered so that occlusions are reduced and more important lines are favored [19, 24]. Indicators for the amount of occlusion in the rendered images are typically based on the "overdraw", e.g., the number of projected line points per pixel [22, 24]. View-dependent opacity control was introduced by Guenther et el. [9] to fade out less important foreground lines. Kanzler et al. [15] use a line cluster hierarchy for view-dependent line density control, by drawing only those lines representative for a cluster and selecting cluster levels according to line importance.

Particle tracing: Particle-based flow exploration relies on interactive frame rates and is, therefore, often concerned with efficient solutions on Graphics Processing Units (GPUs). On recent GPUs it is now possible to interactively trace millions of particles in Cartesian grids using higher order integration schemes [2, 18, 38], and to instantly render these particles using a multitude of different options including oriented point sprites, lines, and more complex geometric representations like bands and tubes [17, 36]. Buerger and co-workers [1] employed GPU particle tracing to efficiently perform probabilistic trajectory tracking and generate particle visitation maps to reveal uncertainties in particle pathways. In our work we perform a similar approach for the instant construction of a particle density map, which is used to visualize the mixing rate of particles from different circulation rolls.

3 Numerical simulation data

The data have been obtained from direct numerical simulations of the 3D Boussinesq equations, which couple the turbulent velocity components u_x, u_y and u_z with the temperature field T. The domain is a flat cuboid with height H, square side length L and a large aspect ratio $\Gamma = L/H$. The bottom (top) plate is held at a constant temperature T_{bottom} (T_{top}) while the sidewalls are adiabatic, i.e. the temperature gradient normal to the wall is zero. All walls are subject to a no-slip boundary condition, i.e. $\mathbf{u} = \mathbf{0}$.

The turbulent flow is then characterised by three dimensionless parameters, the Rayleigh number Ra which quantifies the driving of the turbulence, the Prandtl number Pr which measures the ratio of momentum-to-temperature diffusion of the working fluid and the aspect ratio Γ of the square cell (L/H). For the simulation we employ the spectral element method nek5000, [34] with 2.36 million spectral elements. The parameters of the present data sets are $Ra = 10^5$, $Pr = 0.7$, and $\Gamma = 25$.

Fig. 2 (left) shows that the trajectories in the simulated flow are stochastic and irregular and that it is difficult to detect any relevant coherent structures in this field. However, the time-averaging (Fig. 2 (right)) removes turbulent fluctuations effectively and reveals clearly observable, oriented circulation roll structures which are partly curved. Near the side boundaries, rolls either align with the walls or end perpendicularly. Providing methods to enable a better understanding of these prominent large-scale and long-living structures—so-called superstructures—is the main objective of our work. In particular, we intend to enable an improved analysis of how these superstructures form, how they dominate the global transport of heat, and how they act as barriers to transport.

For temporal and spatial averaging, we estimate a few quantities. A typical particle velocity is given by u_{rms}, the root-mean-square speed computed across the entire domain. The width of the circulation rolls, a, is determined by the first minimum of the spatial correlation functions of vertical velocity component and/or temperature in the midplane [27]. From this, we obtain an estimate of the circulation rolls' circumference, $c \approx \pi(a+H)/2$. The time window for averaging was chosen to be 3τ, where τ is an estimated circulation time of a particle in one of the rolls, $\tau \equiv c/u_{\text{rms}}$. This duration is long enough to smooth out major turbulent fluctuations, but short enough for the superstructures to remain stable.

4 Circulation rolls and defects

To quantitatively analyse the heat exchange between adjacent circulation rolls and the flow behaviour in the vicinity of defects, we first need to extract the rolls' boundaries, and then classify potential defect locations where multiple boundaries meet. Note that even though the rolls stand out clearly when looking from above or below at the flow trajectories, their accurate segmentation is difficult.

4.1 Circulation rolls

To separate the regions of strong circular dynamics in the slowly evolving time-averaged velocity field, it seems promising to detect transport barriers between the forming rolls by means of integration-based measures like the FTLE. The FTLE measures the rate of separation between adjacent particles over a finite time interval with an infinitesimal perturbation in the initial location, and it delineates sharp ridges of high stretching which appear predominantly along the boundaries between adjacent rolls.

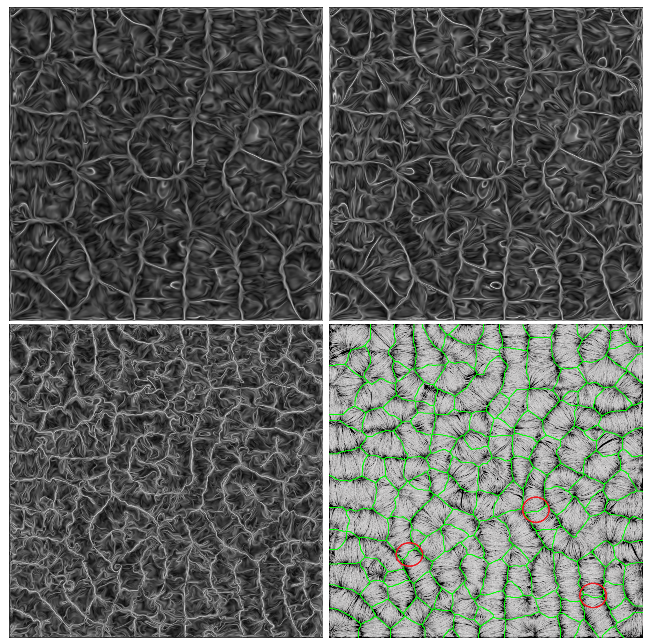

Figure 3: (a-c): FTLE values (using increasing integration time) in the midplane are mapped to brightness. (d): Circulation network extracted in the FTLE field, overlayed on streamline image with selected misclassifications.

We computed a number of FTLE fields with different integration times and step sizes (see Fig. 3(a-c)) to make the values less and more dependent on the amount of local stretching. In the extracted networks, we always observed that some clearly connected roll boundaries split up and artificial separations within single rolls appeared. The most meaningful network we could compute, including hand-tuned modifications like removal of short branches and branch continuation across small gaps [8], is shown in Fig. 3(d). It shows misclassifications and tends to over-segment the domain.

The difficulties in FTLE-based network extraction are due to the specific nature of the used flow fields. Since these fields are time-averages of highly turbulent initial flows, they show small-scale velocity fluctuations which could not be damped out (Fig. 4 (left)). Furthermore, a closer look at regions between adjacent rolls reveals trajectories which cross over to another roll, indicating some heat exchange (Fig. 4 (right)). The FTLE's sensitivity to these local divergences causes the over-segmentation.

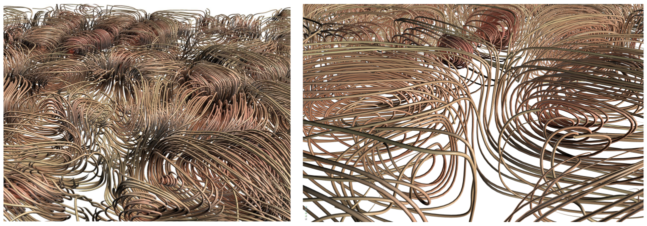

Figure 4: Left: Small-scale turbulent velocity fluctuations remain in the time-averaged flow field. Right: Flow trajectories crossing over to another cell indicate heat transport between rolls.

Our approach to improve the network extraction takes into account the relation between the temperature distribution and the coherent flow structures. Rising (falling) plumes transport warm (cold) fluid from the boundaries into the domain. Hence, we identify the boundaries of circulation rolls as filaments in the average temperature field between z_{bottom} (the heated wall) and z_{top} (the cooled wall). Due to the rising (falling) behaviour of warm (cold) air parcels, the vertical averaging (z-averaging) strengthens in particular the corridors along which the upward (downward) movement occurs. The side-by-side view of a streamline visualization and the horizontally averaged temperature field in a 2D slice, shown in Fig. 5, shows the correspondence between ridges / valleys and roll boundaries. In the temperature plot, the corridors separating the circulation rolls are clearly visible as ridge- and valley-like filaments, i.e., those filaments where the temperature field has a local maximum (minimum) along the direction orthogonal to the filament orientation. The corridors are present already in the unfiltered temperature field in the midplane between the heated bottom and the cooled top wall, yet z-averaging significantly enhances the contrast in values.

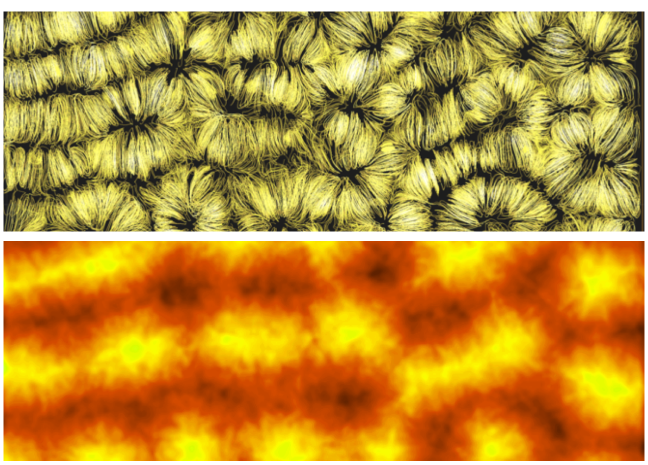

Figure 5: Correspondence between convection rolls (top) and temperature field (bottom, colored from warm (red) to cold (yellow)). Circulation rolls are indicated as spiralling transport structures in the flow. Ridges and valleys in the temperature field indicate boundaries between the rolls.

Based on this observation, the rolls' boundaries are extracted as 1D maxima ridges and minima valleys in the 2D time- and z-averaged temperature field (Fig. 6 (top, left)), i.e., the filaments corresponding to the cores of the plumes along which locally highest and lowest temperature is found. To extract stable structures, the averaged temperature field is first smoothed using a Gaussian filter whose diameter corresponds to the typical width of the circulation rolls, a, described in the previous section.

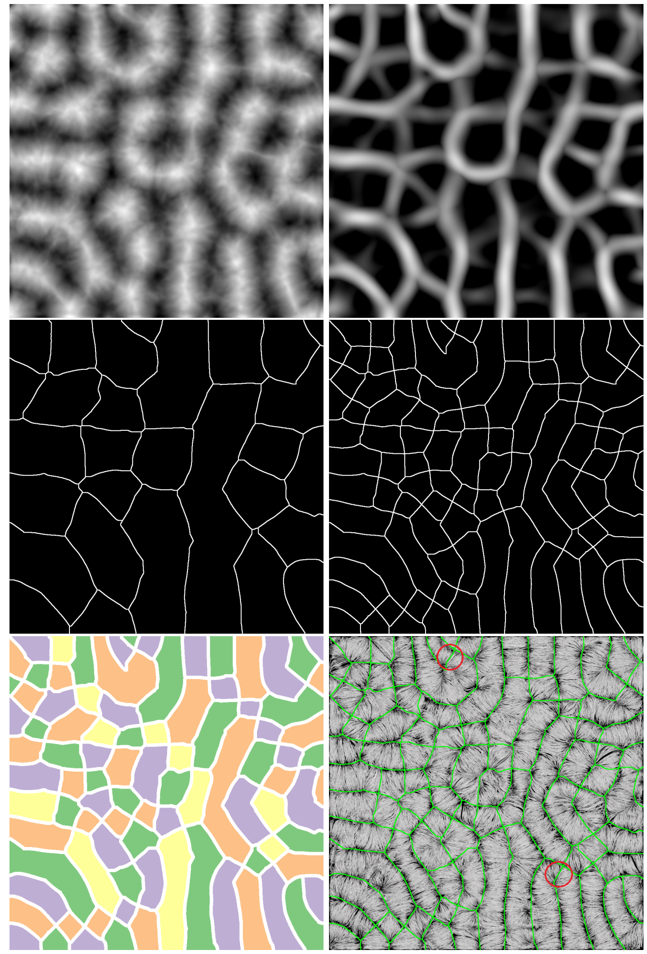

Figure 6: Top: Time- and z-averaged temperature field; main negative eigenvalue of Hessian matrix of smoothed field. Middle: Thinned and pruned ridges; union of ridges and valleys. Bottom: Segmented rolls (with thickened boundaries) using four-coloring scheme; final network overlayed on streamline image with selected misclassifications.

Points on height ridges and valleys can be classified based on geometric properties into the main curvature direction, i.e., the transversal ridge direction [7]. Ridges can be seen as regions of negative curvature and a vanishing gradient into the main curvature direction. Let $T(x,y)$ denote the temperature value at a point (x,y) on the 2D slice $\mathbf{s}$ and let H denote the Hessian matrix of T. We find the negative curvature at a point from the eigenvalues of $H(x,y)$, selecting the lesser value or zero if both eigenvalues are positive (see Fig. 6 (top, right)). Even though the temperature field is smoothed due to temporal and vertical averaging, the ridges that are extracted when points are further classified with respect to the temperature gradient appear fuzzy with significant branching. To avoid this, we perform the following processing steps on the initial 2D eigenvalue field:

After computing the curvature at every point, we first apply a threshold criterion to suppress ridges that are too flat. A binary field is generated, where points at which the curvature is below the selected (negative) threshold are set to 1 (foreground), and 0 (background) otherwise. In the next step, we eliminate spurious features that are caused by small fluctuations near the threshold. Therefore, all orthogonally connected components of foreground points are detected using breadth-first search and points belonging to components which are smaller than a certain threshold (by number of points) are set to 0. We then do the same for small connected components of background points (setting them to 1) to remove small holes in the resulting structures.

To finally reduce the classified foreground points to 1D ridge structures we apply morphological thinning to these points. Thinning is performed by alternating between removing pixels on the north/east boundary and south/west boundary of connected regions of 1-pixels, while preserving the connectivity of the overall shape [10]. The resulting skeleton after thinning is further simplified by pruning the ends of dangling branches, as these do not contribute to our intended segmentation of circulation rolls. We identify end-points of branches as foreground pixels that have only one foreground pixel in their Moore neighbourhood. The corresponding branch is obtained by repeatedly adding the last pixel's foreground neighbour, until that pixel has more than two foreground neighbours itself. All of these branches are deleted from the image by setting these pixels to 0 (Fig. 6 (middle, left)).

Valleys are computed analogously to the ridges in the negated temperature field. As shown in Fig. 6 (middle, right), the union of both skeletons, the so-called heat transport network, segments the domain into a set of convection rolls. For visualization purposes, we let the user decide how many rolls need to be distinguished at once and pick the colors accordingly from a quantitative color scheme at ColorBrewer.org [12]. In Fig. 6 (bottom, left), we show an exemplary four-coloring which ensures that neighbouring circulation rolls are always distinguishable. As the segmentation into circulation rolls is only approximate, we optionally thicken the boundary by several pixels so that pixels right on the boundary are not assigned to either of the two neighbouring rolls.

Fig. 6 (bottom, right) shows the ridge and valley network overlayed on the image of streamlines in the time-averaged flow field. Notably, also the network that is extracted in the temperature field seems to show erroneously classified roll boundaries. Firstly, however, these misclassifications are vastly reduced compared to the FTLE-based classification. Secondly, the majority of these misclassifications could not even be resolved by the domain experts, which were often not sure whether a roll broke up into parts or merged into other rolls. We believe that these uncertain cases can only be resolved by analyzing the flow in the rolls' interiors, by peeling away the outer streamlines.

4.2 Centerline extraction

Besides the swirling motion in the outer parts of the rolls, which can be seen clearly when looking at uniformly seeded trajectories as shown in Fig. 2, not much can be concluded on the flow in a roll's interior. The apparent motion is around a centerline, which is approximately located in the midplane due to the up-down symmetry of the RBC flow. As we will demonstrate in Sec. 5, extracting streamlines close to these centerlines helps analysing how the flow behavior changes in the center of the rolls, hinting toward non-uniform and roll-specific rotational circulations. We propose to first extract approximate centerlines from the ridge/valley network, and then compute the closeness of streamlines to these approximate centerlines. The closeness values are then considered in the visualization process to thin out the outer rotating lines.

From the binary field in which ridge/valley pixels are classified, we compute a distance field where every pixel stores the shortest distance to any of the ridge/valley pixels (Fig. 7 (left)). In this field, the centerlines appear as maxima ridges, giving rise to the following computation of the closeness of a streamline to a centerline: We traverse along all streamlines and retrieve at every point a distance value by projecting that point along the z-axis into the distance field (i.e. by setting its z coordinate to the z location of the midplane). From this value and the distance of the initial streamline point to the midplane, the distance to the closest centerline is computed by the Pythagorean theorem. We do this for every point along a streamline, and compute the root mean square (RMS) over all computed distances. Streamlines having a RMS below a selected threshold are then classified as close to a centerline. Fig. 7 (right) shows all streamlines, projected and additively blended into the midplane using a blend factor that is inversely proportional to closeness.

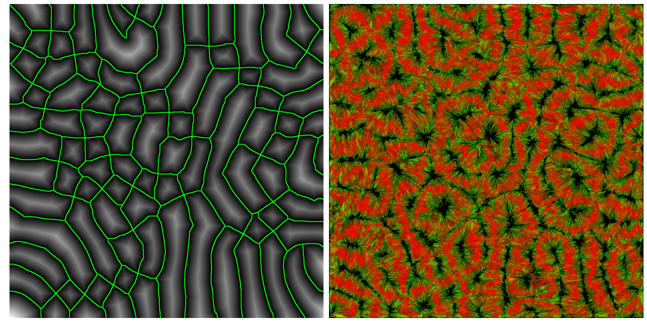

Figure 7: Left: Distance values to roll boundaries (green), color coded from black (small values) to white (large values). Right: All 3D trajectories are colored from green (high distance to midplane) to red (short distance to midplane) and projected into the midplane using accumulative blending.

4.3 Defects

In the streamlines of the time-averaged velocity field, one can observe defects where a third circulation roll extends into a regular convection pattern formed by two rolls (marked with circles in Fig. 3 (right)). In the plumes' cores, such defects are indicated by so-called trisectors where three network edges meet. Analysing the flow in the surrounding of these locations can be performed interactively, by visualizing the extracted network and seeding particles close to the shown defect locations. However, to shed light on the similarities and differences between the flow in the vicinity of different defects, an automated structural or statistical analysis and comparison of the flow in these regions is necessary. Thus, we provide a method for detecting potential defects automatically, which can be used to simplify user interaction, by automatically centering a probe at such locations, or for a comparative defect analysis.

In a first attempt to detect trisectors, we used the tensor method proposed by Delmarcelle [6] to quantify the local alignment of the circulation rolls. The method is based on the observation that for linear interpolation only two types of isolated degenerate points in a tensor field exist, so-called wedge points and trisectors, with a winding number (or tensor index) of $1/2$ and $-1/2$, respectively. The winding number can be computed by integrating over the angle between the eigenvectors of the derivative tensor along an enclosing Jordan curve.

When applying the tensor method to the temperature field, however, we noticed that found trisectors were detected a significant distance away from the points where the circulation rolls actually meet. We attribute this fact to the smoothing of the temperature field, which needs to be performed in order to obtain stable tensor components and which smears out the defect locations within the applied filter region. Without smoothing, on the other hand, many false positives are detected which do not correspond to any defects.

To overcome this limitation we extract the trisectors directly from the network structure. We use points where three ridges meet as potential candidates for trisectors. These are defined as pixels which lie on a ridge and whose Moore neighbourhood contains exactly three contiguous runs of ridge pixels (an example is shown in Fig. 8 (left)). To detect these, we compute the eight differences, Δ_i, between adjacent pixels in the Moore neighbour and check that $\sum_i |\Delta_i| = 6$.

Occasionally, the proposed algorithm yields trisectors in nearby pairs. These correspond to four ridges approximately meeting at a

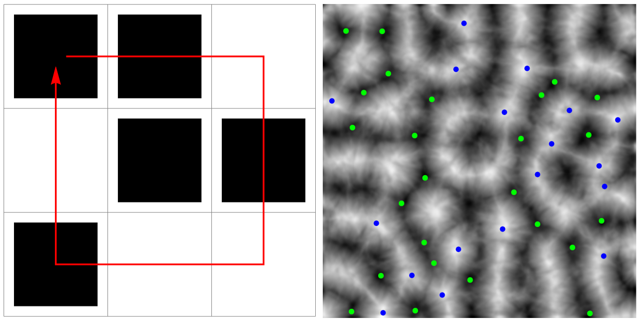

Figure 8: Left: Trisectors are detected as ridge pixels whose Moore neighbourhood contains exactly three contiguous runs of ridge pixels. Right: Extracted trisectors after filtering out nearby pairs. Blue (green) trisectors were found in the ridge (valley) field.

point, not just three. Therefore, we consider them false positives and reject all trisectors which appear within a certain distance from another trisector. Fig. 8 (right) shows the resulting set of trisectors.

5 Trajectory-based transport visualization

As shown in Fig. 2, by seeding streamlines uniformly in the entire domain, the dominant roll-like heat transport patterns can be observed clearly. When looked from the outside, the visualization seems to indicate homogeneous flow patterns in the interior of circulation rolls. To investigate these patterns in more detail, we provide means to investigate the circulation in the rolls' interior via density-controlled streamline rendering and 3D particle visualization.

5.1 Density-controlled streamline rendering

While it is possible to manually select seeding locations in a roll's interior and, thus, focus solely on interior trajectories, our partners from computational fluid dynamics are especially interested in comparing the characteristics of interior and outer trajectories, for instance, their merging and crossing behaviour. Therefore, we decided to use line rendering via density-control in screen-space as proposed by Kanzler et al. [15], to automatically reduce the density of outer streamlines yet preserving some of them as context information.

Line rendering via density-control in screen-space uses a hierarchy of line clusters to thin out less important lines in a view-dependent way. It assigns an importance value to every segment along a line, and combines this measure with view-dependent measures like per-pixel overdraw to assign to each segment a level in the cluster hierarchy. That means, that with decreasing importance the segment is assigned to ever coarser levels. A segment is only rendered if at the assigned level it is the representative line for the cluster it belongs to, meaning that line segments are thinned out with decreasing importance or when occluding more important ones. For the construction of a line cluster hierarchy we follow previous work and use a variant of Agglomerative Hierarchical Clustering (AHC) with single linkage and graph matching to construct a fully balanced cluster tree. Streamlines are extracted from the time-averaged flow field in a pre-process, and they are rendered on the GPU using the geometry shader to construct tubes with user-selected diameter.

Fig. 9 shows the use of line curvature, which is commonly applied in screen-space approaches, as importance measure. Surprisingly, the flow in the rolls interior seems far from homogeneous and perfectly circular. Especially toward a roll's centerline the trajectories often become helix-like with increasing elongation along the centerline. Unfortunately, with increasing helix-likeliness, a line's curvature decreases so that curvature-based importance cannot emphasize such trajectories. To circumvent this behaviour, we propose to combine the curvature-based criterion with a criterion based on the closeness of trajectories to the rolls' centerlines.

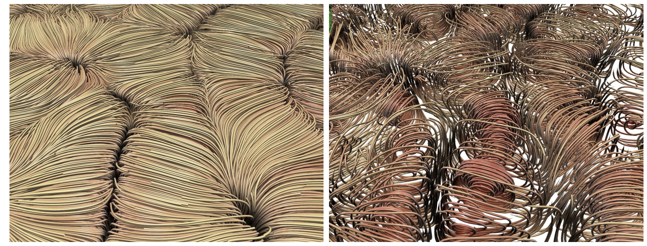

Figure 9: Left: Outer trajectories occlude interior ones. Right: Line-density control based on line curvature thins out outer trajectories and reveals the interior flow structure in a superstructure circulation roll.

As discussed in Sec. 4, for every segment of a trajectory we can compute a distance value indicating the shortest distance to any of the centerlines in the midplane. By using an importance measure that blends linearly between curvature and closeness to centerline, we can smoothly fade between highly-curved roll-like and elongated helix-like circulation patterns. In this way, the visualization provides a context view via thinned-out outer trajectories, and a focus view comprised of the interior curved trajectories and the trajectories close to the rolls' centerlines. Fig. 10 demonstrates the effectiveness of this approach for revealing the interior structure of circulation rolls.

Figure 10: Focus and context streamline visualization: Few outer trajectories indicate the shape and size of circulation rolls, highly-curved roll-like interior lines and elongated helix-like lines close to the rolls' centerlines indicate specific circulation patterns.

5.2 Transport analysis using 3D particle visualization

While streamline visualization can reveal the structure of the flow, it cannot show the amount and speed of mixing between particles which were initially seeded in different rolls. To analyse such behaviour, we use the circulation network to steer the seeding of particles in combination with interactive 3D particle visualization to analyse their trajectories and mixing behaviour. The techniques have been integrated into a GPU-based particle visualization tool, and particles can either be seeded once before the advection starts, or continually while the advection is running.

Particles are rendered as small transparent ellipses with an aspect ratio and orientation that is continually adapted in a GPU geometry shader to indicate the current flow direction. To allow for the correct blending of particles when they are made semi-transparent, they are depth-sorted in every frame on the GPU using NVIDIA's radix sort implementation in the CUB high-performance library for CUDA kernel programming. To provide context, an opaque or transparent 2D slice—texture mapped with the extracted ridge pattern—can be visualized simultaneously.

5.3 Segmentation-based seeding and coloring

To focus the analysis on the heat transport between certain circulation rolls, rolls can be selected interactively by picking the corresponding regions in the segmentation. Particle seeding is then restricted to these rolls via rejection sampling in the segmented field. This enables the user to investigate along which pathways the heat of specific circulation rolls is transported through the domain and, in particular, into which adjacent rolls heat is transported. By seeding particles densely, the user can analyse quantitatively how many particles stay in their initial circulation roll or leave this roll, and where in the domain the transition from one roll to another occurs.

By assigning each particle the color of the circulation roll in which it was seeded initially, the mixing of fluid from different rolls can be visualized. The coloring is based on the assignment of colors to rolls as shown in Fig. 6 (bottom, first), and it establishes an effective visual connection between particles and rolls. In Fig. 11, two rolls were selected as seed areas for particles, which were then traced over a certain time period. The different degrees of mixing and the mixing patterns are revealed clearly. Besides showing between which rolls the mixing is low or high, one can also analyse quantitatively how mixing proceeds within a certain roll, i.e. whether mixing affects the entire roll or only particular streams within one roll. It was an interesting immediate finding that the mixing behavior strongly differs between rolls. Notably, also the speed of mixing can be monitored when using interactive particle tracing, which cannot easily be achieved using streamline visualization.

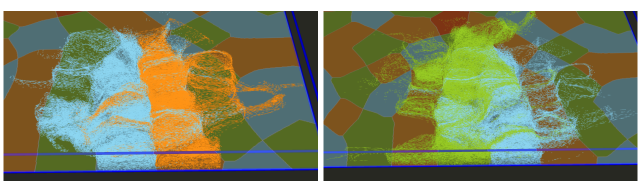

Figure 11: Weak (left) and strong (right) mixing of air between adjacent rolls is revealed by coloring particles according to the roll in which they were seeded initially. While mixing proceeds along specific streams on the left, it evolves across the entire rolls on the right.

The user can also focus the analysis on the flow in the surrounding of a defect, by letting the system restrict the seeded particles to a box around the defect location or the three circulation rolls which touch a trisector (see Fig. 12). Such an investigation can also be performed via streamlines, yet by using particles the speed of movement and mixing can be perceived, too. Around the selected defect, a specific phenomenon with respect to the predominant mixing direction could be observed: The mixing of particles from one roll into another one predominantly occurs between the blue and green roll (inside the green roll), whereas matter from the orange roll is transported entirely out of this three-roll system, and no matter enters the green or orange roll. Interestingly, and unknown before, this anisotropic mixing behavior, indicative of some specific time-averaged mass exchange at the defect location, does not only occur around one particular trisector but can also be found at several other locations.

5.4 Dwell time computation

We introduce the dwell time as an indicator of how much time a particle has spent in its current roll. It is computed by keeping track of when each particle has last crossed a roll boundary. As there is a certain amount of uncertainty in the position of the roll boundaries, we require a particle to spend an (adjustable) minimum amount of time in a new roll before the particle is considered to have crossed

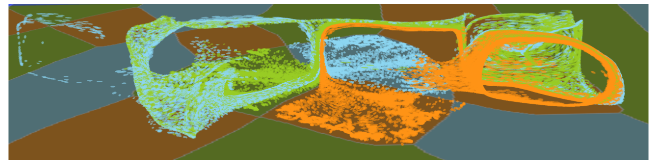

Figure 12: Defect visualization: Highly anisotropic mixing behaviour between neighbouring rolls occurs in the vicinity of potential defectany of the extracted defect locations.

over. Thus, we do not reset the dwell time if a particle just barely crosses a boundary and then returns to its previous circulation roll. Once a particle crosses over a boundary for the required minimum duration, its dwell time is reset to that duration. Additionally, we also keep track of the circulation roll in which the particle has stayed for the longest time.

Visualizing the per-particle dwell time can be performed in the following ways to provide specific insights into the transport and mixing behaviour of the flow:

1. By highlighting particles with a large dwell time, the user can focus on those particles that form stable structures which remain within one circulation roll for a long time (see Fig. 13 (top)).

2. In contrast, by highlighting particles with a small dwell time, particles which have only recently moved from one circulation roll to another are emphasized. This guides toward those regions where particles are exchanged along a given roll boundary, if at all (see Fig. 13 (bottom)).

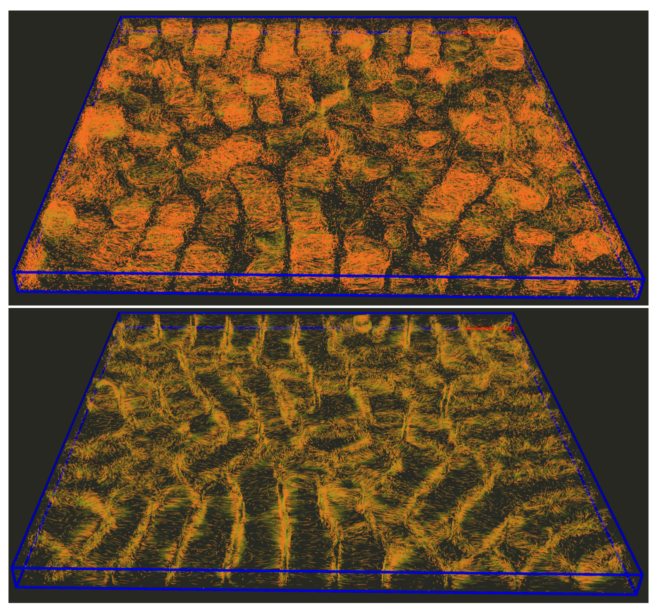

Figure 13: The opacity of particles corresponds to dwell time, to highlight stable structures (top; long dwell time) or particle exchange between adjacent rolls (bottom; short dwell time).

When using the dwell time to analyse the mixing behaviour it is important to consider that the threshold indicating that a particle has crossed over a boundary, i.e. the particle has stayed in a roll for a sufficient amount of time, is not too large. Assume a particle circulates quickly inside one of the rolls, but because of the uncertainty in the exact boundary location it starts in another roll and crosses over to this roll every rotation again frequently. If the threshold was chosen too large, this particle will never be recognised as belonging to its true circulation roll, but stays in the starting roll even if the total time spent there is much less than it spent in the true roll.

5.5 Density maps

While interactive particle visualization can show very clearly where mixing of fluid from different rolls occurs, it is difficult to perform a quantitative analysis of the mixing ratios with this approach. Therefore, we also visualize the mixing ratio directly, as a scalar quantity that is computed instantly during particle integration. To obtain the mixing ratio we first compute a separate density map for each circulation roll of interest—interactively selected by the user—as a domain-filling grid. The resolution of the grid is slightly coarser than the resolution of the grid at which the simulation data is stored. For each cell in the grid, we count how many particles seeded in that roll currently lie inside the cell.

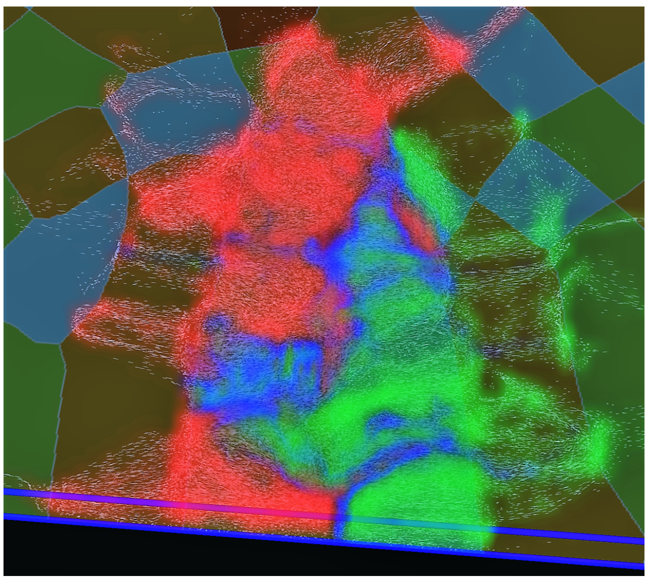

Figure 14: Density map visualization of mixing ratio of particles from two adjacent circulation rolls. Mixing is approximately equal in blue regions.

The mixing ratio between two selected rolls is then computed from the corresponding density maps in each grid cell. Given the numbers of particles in a cell from rolls i and j as N_i and N_j, the mixing ratio is defined as $\rho = N_i/(N_i + N_j)$. This gives a value between 0 and 1, with 0 or 1, respectively, indicating only particles from rolls i or j, and 0.5 indicating an equal number of particles from both rolls.

To visualize the mixing ratio, direct volume rendering is used. We map the ratio to color using an adjustable transfer function, where opacity is determined by the sum of the particle counts, $N_i + N_j$, normalized based on the maximal value across the entire domain. This makes it easy to highlight separately the regions where particles from either circulation roll dominate, or where they are present in roughly equal quantities. The visualization of mixing ratios is shown in Fig. 14, where the coloring separates regions of weak and strong mixing. Alternatively, Fig. 15 shows the same region with an isosurface at a mixing ratio of $\rho = 0.5$ at two different times during the simulation. The isosurface is only computed in regions where $N_i + N_j > 0$, so that the mixing ratio is well defined, and opacity is mapped as described.

The visualization of the mixing ratios turns out to be a powerful option when used interactively. At the same time it enables to detect locations where leakage between two adjacent rolls occurs and whether this leakage is from both sides of the boundary or only unidirectional. Furthermore, by watching the changing mixing ratios over time, it becomes immediately clear at which speed the leakage occurs in different regions. The possibility to visualize an isosurface in the density field, and in particular its evolution over time, facilitates a quantitative analysis of the mixing behaviour between different rolls. In the future, we intend to consider the extension of this functionality to enable the analysis of mixing ratios between more than two rolls simultaneously.

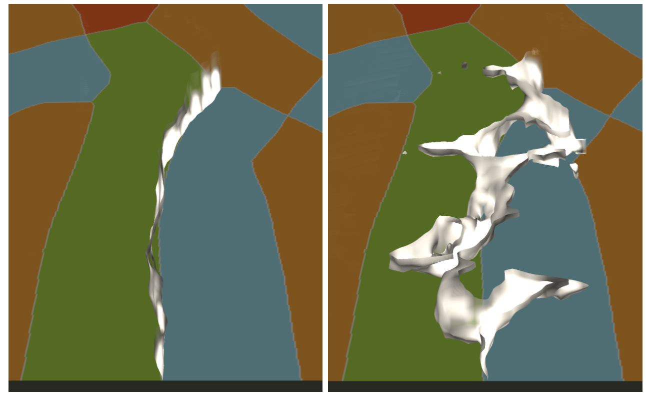

Figure 15: Isosurface indicating a 50% mixing ratio between the green and blue circulation rolls. The lefthand snapshot was taken shortly after seeding the particles in their respective cells, the righthand snapshot after letting the particles mix for a while.

6 Concluding Discussion and Future Work

In this work we have introduced visualization techniques to enable further progress in the understanding of convective heat transport. These techniques build upon the segmentation of circulation rolls using ridge extraction in the z-averaged temperature field and make use of the segmentation to focus the analysis on specific transport patterns. The segmentation is tailored to the up-down symmetry of the RBC flow, which forms corridors of falling and rising heat patterns. In a follow-up application study, we aim to apply this segmentation to RBC flows with different Prandtl and Rayleigh numbers to shed light on the forming of superstructures in different turbulent regimes.

Building upon the segmentation, we introduced an importance measure to reveal the internal structure of the circulation rolls via streamline visualization using density control in screen-space. Interactive particle visualization enables the user to investigate the speed of transport and the mixing ratio between particles released in different rolls. By using these visualization options, major distinct trends in the heat transfer within and across circulation rolls can be revealed.

On our target machine, an Intel Xeon E5-1650 v3 CPU with 6×3.50 GHz, 32 GB RAM, and an NVIDIA GeForce GTX 970 graphics card with 4 GB VRAM, the pre-process to compute and cluster trajectories of 10.000 particles using 200 integration steps takes roughly 1 minute, using a fourth order Runge-Kutta integrator and Agglomerative Hierarchical Clustering with a mean-of-closest-point distance. Particle visualization can be performed in a fully interactive way: One integration step of 100,000 particles using a fourth order Runge-Kutta integrator takes about 10 milliseconds. Segmentation of the temperature field to extract the circulation network and trisectors is performed on the CPU. It takes roughly 1 second for one slice of the data on a 1024×1024 Cartesian grid. The entire dataset has a size of $1024 \times 1024 \times 32$.

The most important findings of our analysis are as follows: Firstly, by means of density-controlled streamline rendering using the closeness to the rolls' centerlines, we could reveal the different types of flow patterns in the interior of circulation rolls. Fig. 16 shows in particular the existence of rather elongated helix-like streamlines close to a roll's centerline. The more circular patterns occur with vastly different frequency.

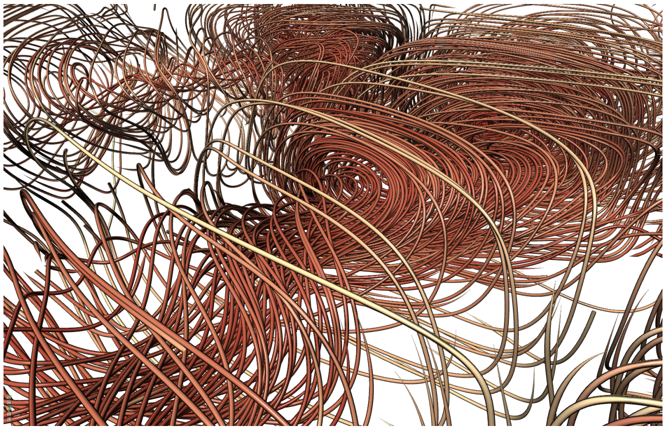

Figure 16: Elongated helix-like streamlines are found in the rolls' cores.

Secondly, we could effectively reveal the different mixing characteristics between different rolls. As it is also shown in Fig. 17, different rolls exhibit different mixing strengths, and even spatially varying mixing behavior. One important aspect our partners from computational fluid dynamics want to investigate in the future is whether there are relationships between the mixing behavior, the rolls' lengths, and the rolls' locations.

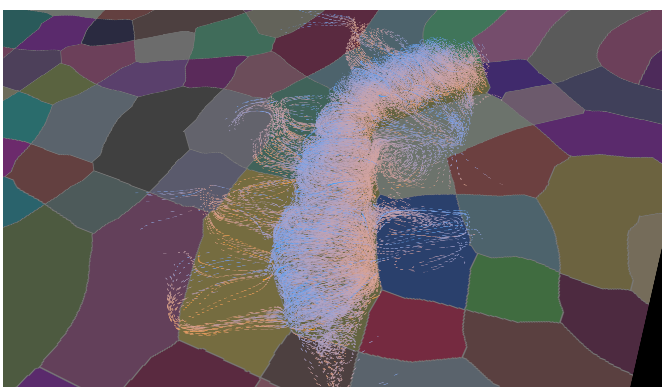

Figure 17: Particle visualization (temperature mapped to color) reveals rolls with vastly different and inhomogeneous mixing behaviour.

Thirdly, it is surprising that very similar mixing behaviours can be observed around defects at different locations in in the domain. Fig. 18 shows mixing patterns close to a defect which look very similar to the pattern found in the vicinity of the defect visualized in Fig 12. A more detailed and probably statistical evaluation of these patterns can help to better understand the origin of defects, and the specific constellation in which they emerge.

Our techniques enable an interactive exploration of the major heat transport patterns in 3D time-averaged RBC flow. However, the temporal evolution and the multiscale nature of superstructures has not been addressed so far. This requires to perform roll extraction not only in one single time-averaged field, but in a sequence of time-averaged evolutions, including the tracking of structures over

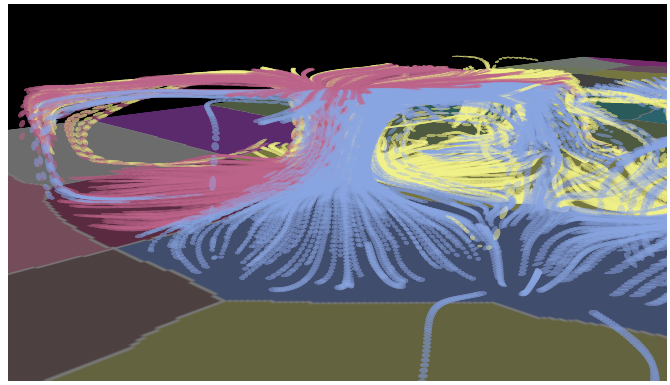

Figure 18: Similar mixing behaviour in the vicinity of defects at different locations.

time to analyse their dynamical behaviour. Furthermore, it needs to be shed light on the effect of different averaging periods on the superstructures in the resulting fields. Therefore, we aim to analyse when and where structures appear when multiscale averaging operators are used. The interplay between structures across scale and time, and their relation to specific patterns in the particle movements at different scales need to be further investigated. By considering the emergence of superstructures over multiple scales, we hope to enable an automated segmentation of circulation rolls. Furthermore, we aim to shed light on the origin and evolution of defects in the instationary fields, by analysing when and in which concrete settings they tend to emerge and for how long they remain alive.

7 Acknowledgments

This work has been done within the Priority Programme Turbulent Superstructures (SPP 1881), funded by the German Research Foundation (DFG). Computing resources were provided by the Large Scale Project praise of the Gauss Centre for Supercomputing.

References

[1] K. Bürger, R. Fraedrich, D. Merhof, and R. Westermann. Instant visitation maps for interactive visualization of uncertain particle trajectories. *Proc. SPIE*, 8294:8294–8294, 2012. doi: 10.1117/12.906872

[2] K. Bürger, J. Schneider, P. Kondratieva, J. Krüger, and R. Westermann. Interactive visual exploration of instationary 3D-flows. pp. 251–258, 2007.

[3] F. Busse. The twins of turbulence research. *Physcs Online Journal*, 5:4–, 01 2012. doi: 10.1103/Physics.5.4

[4] F. Chillà and J. Schumacher. New perspectives in turbulent Rayleigh-Bénard convection. *The European Physical Journal E*, 35(7):1–25, 2012.

[5] M. Cross and H. Greenside. *Pattern Formation and Dynamics in Nonequilibrium Systems*. Cambridge University Press, 2009.

[6] T. Delmarcelle. The visualization of second-order tensor fields. 1995.

[7] D. Eberly, R. Gardner, B. Morse, S. Pizer, and C. Scharlach. Ridges for image analysis. *J. Math. Imaging Vis.*, 4(4):353–373, 1994. doi: 10.1007/BF01262402

[8] F. Ferstl, K. Bürger, H. Theisel, and R. Westermann. Interactive separating streak surfaces. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1569–1577, 2010. doi: 10.1109/TVCG.2010.169

[9] T. Günther, C. Rössl, and H. Theisel. Opacity optimization for 3D line fields. *Proc. ACM SIGGRAPH*, 32(4):120, 2013.

[10] Z. Guo and R. W. Hall. Parallel thinning with two-subiteration algorithms. *Communications of the ACM*, 32(3):359–373, 1989.

[11] G. Haller. Distinguished material surfaces and coherent structures in three-dimensional fluid flows. *Physica D: Nonlinear Phenomena*, 149(4):248 – 277, 2001. doi: 10.1016/S0167-2789(00)00199-8

[12] M. Harrower and C. A. Brewer. Colorbrewer.org: an online tool for selecting colour schemes for maps. *The Cartographic Journal*, 40(1):27–37, 2003.

[13] H. Hochstetter, M. Wurm, and A. Kolb. Vector field visualization of advective-diffusive flows. *Comput. Graph. Forum*, 34(3):481–490, 2015.

[14] B. Jobard and W. Lefer. Creating evenly-spaced streamlines of arbitrary density. In *Proc. Eurographics Workshop on Visualization in Scientific Computing*, pp. 43–55, 1997.

[15] M. Kanzler, F. Ferstl, and R. Westermann. Line density control in screen-space via balanced line hierarchies. *Computers & Graphics*, 61:29 – 39, 2016. doi: 10.1016/j.cag.2016.08.001

[16] G. K. Karch, F. Sadlo, D. Weiskopf, C.-D. Munz, and T. Ertl. Visualization of advection-diffusion in unsteady fluid flow. *Computer Graphics Forum*, 31:1105–1114, 2012.

[17] P. Kondratieva, J. Kruger, and R. Westermann. The application of gpu particle tracing to diffusion tensor field visualization. In *VIS 05. IEEE Visualization, 2005.*, pp. 73–78, 2005. doi: 10.1109/VISUAL.2005.1532780

[18] J. Krüger, P. Kipfer, P. Kondratieva, and R. Westermann. A particle system for interactive visualization of 3d flows. *IEEE Transactions on Visualization and Computer Graphics*, 11(6):744–756, 2005.

[19] T.-Y. Lee, O. Mishchenko, H.-W. Shen, and R. Crawfis. View point evaluation and streamline filtering for flow visualization. In *Proc. IEEE PacificVis*, pp. 83–90, 2011. doi: 10.1109/PACIFICVIS.2011.5742376

[20] L. Li, H.-H. Hsieh, and H.-W. Shen. Illustrative streamline placement and visualization. In *Proc. IEEE PacificVis*, pp. 79–86, 2008. doi: 10.1109/PACIFICVIS.2008.4475462

[21] Z. Liu, R. Moorhead, and J. Groner. An advanced evenly-spaced streamline placement algorithm. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):965–972, 2006. doi: 10.1109/TVCG.2006.116

[22] J. Ma, C. Wang, and C.-K. Shene. Coherent view-dependent streamline selection for importance-driven flow visualization. *Proc. SPIE 8654, Visualization and Data Analysis*, 2013. doi: 10.1117/12.2001887

[23] G. M. Machado, S. Boblest, T. Ertl, and F. Sadlo. Space-time bifurcation lines for extraction of 2D Lagrangian coherent structures. *Computer Graphics Forum*, 35(3):91–100, 2016. doi: 10.1111/cgf.12885

[24] S. Marchesin, C.-K. Chen, C. Ho, and K.-L. Ma. View-dependent streamlines for 3D vector fields. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1578–1586, 2010. doi: 10.1109/TVCG.2010.212

[25] O. Mattausch, T. Theußl, H. Hauser, and E. Gröller. Strategies for interactive exploration of 3D flow using evenly-spaced illuminated streamlines. In *Proceedings of the 19th Spring Conference on Computer Graphics*, pp. 213–222, 2003. doi: 10.1145/984952.984987

[26] T. McLoughlin, M. Jones, R. Laramee, R. Malki, I. Masters, and C. Hansen. Similarity measures for enhancing interactive streamline seeding. *IEEE Transactions on Visualization and Computer Graphics*, 19(8):1342–1353, 2012. doi: 10.1109/TVCG.2012.150

[27] A. Pandey, J. D. Scheel, and J. Schumacher. Turbulent superstructures in rayleigh-bénard convection. *Nature Communications*, 9, 2018(art. 2118), Jun 2018. doi: 10.1038/s41467-018-04478-0

[28] R. Peikert and F. Sadlo. Visualization methods for vortex rings and vortex breakdown bubbles. In *Proceedings of the 9th Joint Eurographics / IEEE VGTC Conference on Visualization*, pp. 211–218, 2007. doi: 10.2312/VisSym/EuroVis07/211-218

[29] R. Peikert and F. Sadlo. Height ridge computation and filtering for visualization. In *2008 IEEE Pacific Visualization Symposium*, pp. 119–126, 2008. doi: 10.1109/PACIFICVIS.2008.4475467

[30] F. Sadlo, G. Karch, and T. Ertl. Topological features in time-dependent advection-diffusion flow. In *Topological Methods in Data Analysis and Visualization III*, pp. 217–231, 2014.

[31] F. Sadlo and R. Peikert. Efficient visualization of lagrangian coherent structures by filtered AMR ridge extraction. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1456–1463, 2007. doi: 10.1109/TVCG.2007.70554

[32] F. Sadlo and D. Weiskopf. Time-dependent 2-d vector field topology: An approach inspired by lagrangian coherent structures. *Computer Graphics Forum*, 29(1):88–100, 2010. doi: 10.1111/j.1467-8659.2009.01546.x

[33] 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. doi: 10.1109/TVCG.2010.133

[34] J. D. Scheel, M. S. Emran, and J. Schumacher. Resolving the fine-scale structure in turbulent Rayleigh–Bénard convection. *New Journal of Physics*, 15(11):113063, 2013.

[35] B. Schindler, R. Fuchs, S. Barp, J. Waser, A. Pobitzer, R. Carnecky, K. Matkovic, and R. Peikert. Lagrangian coherent structures for design analysis of revolving doors. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2159–2168, 2012. doi: 10.1109/TVCG.2012.243

[36] M. Schirski, T. Kuhlen, M. Hopp, P. Adomeit, S. Pischinger, and C. Bischof. Efficient visualization of large amounts of particle trajectories in virtual environments using virtual tubelets. In *VRCAI '04: Proceedings of the 2004 ACM SIGGRAPH international conference on Virtual Reality continuum and its applications in industry*, pp. 141–147, 2004.

[37] J. Schumacher. Lagrangian dispersion and heat transport in convective turbulence. *Physical Review Letters*, 100:134502, 2008.

[38] H.-W. Shen, G.-S. Li, and U. D. Bordoloi. Interactive visualization of three-dimensional vector fields with flexible appearance control. *IEEE Transactions on Visualization and Computer Graphics*, 10(4):434–445, 2004.

[39] O. Shishkina and C. Wagner. Analysis of sheet-like thermal plumes in turbulent Rayleigh-Bénard convection. *Journal of Fluid Mechanics*, 599:383–404, 2008.

[40] D. Stalling, M. Zockler, and H.-C. Hege. Fast display of illuminated field lines. *IEEE Transactions on Visualization and Computer Graphics*, 3(2):118–128, 1997. doi: 10.1109/2945.597795

[41] H. Theisel, T. Weinkauf, H. C. Hege, and H. P. Seidel. Saddle connectors - an approach to visualizing the topological skeleton of complex 3d vector fields. In *IEEE Visualization, 2003. VIS 2003.*, pp. 225–232, 2003. doi: 10.1109/VISUAL.2003.1250376

[42] H. Theisel, T. Weinkauf, H. C. Hege, and H. P. Seidel. Stream line and path line oriented topology for 2d time-dependent vector fields. In *IEEE Visualization 2004*, pp. 321–328, 2004. doi: 10.1109/VISUAL.2004.99

[43] G. Turk and D. Banks. Image-guided streamline placement. In *Proc. ACM SIGGRAPH*, pp. 453–460, 1996. doi: 10.1145/237170.237285

[44] M. Uffinger, F. Sadlo, and T. Ertl. A time-dependent vector field topology based on streak surfaces. *IEEE Transactions on Visualization and Computer Graphics*, 19(3):379–392, 2013. doi: 10.1109/TVCG.2012.131

[45] V. Verma, D. Kao, and A. Pang. A flow-guided streamline seeding strategy. In *Proc. IEEE Visualization*, pp. 163–170, 2000.

[46] T. Weinkauf, H. Theisel, H.-C. Hege, and H.-P. Seidel. Boundary switch connectors for topological visualization of complex 3d vector fields. In *Proceedings of the Sixth Joint Eurographics - IEEE TCVG Conference on Visualization*, VISSYM'04, pp. 183–192, 2004. doi: 10.2312/VisSym/VisSym04/183-192

[47] L. Xu, T.-Y. Lee, and H.-W. Shen. An information-theoretic framework for flow visualization. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1216–1224, 2010. doi: 10.1109/TVCG.2010.131

[48] X. Ye, D. Kao, and A. Pang. Strategy for seeding 3D streamlines. In *Proc. IEEE Visualization 2005*, pp. 471–478, 2005. doi: 10.1109/VISUAL.2005.1532831

[49] H. Yu, C. Wang, C.-K. Shene, and J. H. Chen. Hierarchical streamline bundles. *IEEE Transactions on Visualization and Computer Graphics*, 18(8):1353–1367, 2012. doi: 10.1109/TVCG.2011.155

[50] Q. Zhou, C. Sun, and K.-Q. Xia. Morphological evolution of thermal plumes in turbulent Rayleigh-Bénard convection. *Physical Review Letters*, 98:074501, 2007.