

Interactive Virtual Probing of 4D MRI Blood-Flow

Roy van Pelt, Javier Oliván Bescós, Marcel Breeuwer, Rachel E. Clough, M. Eduard Gröller, *Member, IEEE*, Bart ter Haar Romeny, and Anna Vilanova, *Member, IEEE*




Fig. 1. Interactive virtual probe with flow visualization approaches, enabling exploration of cardiovascular 4D MRI blood-flow data. Color in the leftmost rendition encodes the blood-flow vorticity, while color in other renditions conveys the local blood-flow speed.

Abstract—Better understanding of hemodynamics conceivably leads to improved diagnosis and prognosis of cardiovascular diseases. Therefore, an elaborate analysis of the blood-flow in heart and thoracic arteries is essential. Contemporary MRI techniques enable acquisition of quantitative time-resolved flow information, resulting in 4D velocity fields that capture the blood-flow behavior. Visual exploration of these fields provides comprehensive insight into the unsteady blood-flow behavior, and precedes a quantitative analysis of additional blood-flow parameters. The complete inspection requires accurate segmentation of anatomical structures, encompassing a time-consuming and hard-to-automate process, especially for malformed morphologies. We present a way to avoid the laborious segmentation process in case of qualitative inspection, by introducing an interactive virtual probe. This probe is positioned semi-automatically within the blood-flow field, and serves as a navigational object for visual exploration. The difficult task of determining position and orientation along the view-direction is automated by a fitting approach, aligning the probe with the orientations of the velocity field. The aligned probe provides an interactive seeding basis for various flow visualization approaches. We demonstrate illustration-inspired particles, integral lines and integral surfaces, conveying distinct characteristics of the unsteady blood-flow. Lastly, we present the results of an evaluation with domain experts, valuing the practical use of our probe and flow visualization techniques.

Index Terms—Probing, Flow visualization, Illustrative visualization, Multivalued images, Phase-contrast cine MRI.

1 INTRODUCTION

Morphology of cardiovascular tissue is significantly influenced by the unsteady behavior of flowing blood, and hence pathogenesis of cardiovascular diseases (CVD) is directly affected by the blood-flow dynamics. CVD comprises a group of conditions with a prevalence of over thirty percent of the American population, and is currently the leading cause of death worldwide [29].

Diagnosis and prognosis of CVD is largely based on acquired morphology, possibly extended by functional information from cine scans. In clinical practice, the influence of blood-flow dynamics is largely ne-

glected. Strong clinical trials are necessary to understand the measured hemodynamics, promising great diagnostic and prognostic value.

This especially holds true for complex cardiovascular conditions, causing malformations of the anatomical structures. In particular, congenital heart diseases (CHD) imply anomalous hemodynamics that strongly influence the progression of the innate defects. For the adult case, a noteworthy application is the aortic dissection, which is caused by a tear in the inner aortic wall. This allows blood to flow between the disintegrated layers of the vessel wall, resulting in a high risk of rupture. Again, the blood-flow behavior plays a predominant role in the course of the condition.

Blood-flow information is mostly obtained by quantitative ultrasound (US) acquisition. US is a cost-effective modality, providing high temporal resolution. However, US acquisition requires a skilled operator, and is generally subject to a substantial amount of noise. Consequently, this modality is less suitable for challenging cardiovascular conditions. Alternatively, computed tomography (CT) provides a limited number of blood-flow acquisition sequences, delivering better signal-to-noise ratios. CT has the drawback of exposing the patient to harmful radiation, which is impermissible for young CHD patients. Instead, we focus on non-invasive Phase-Contrast MRI acquisition, providing 3D quantitative measurements of blood-flow velocities throughout the cardiac cycle. A representative subset of the data is depicted in figure 2.

The analysis of hemodynamic characteristics involves qualitative and quantitative inspection of the blood-flow field. Physicians investigate both the spatiotemporal flow behavior, as well as derived measures, such as the mean flux or cardiac output. Currently, the workflow for qualitative and quantitative analysis commences with an elaborate segmentation of the anatomical structures of interest. This is generally

• Roy van Pelt, Marcel Breeuwer, Bart ter Haar Romeny and Anna Vilanova are with the department of Biomedical Engineering, within the group of Biomedical Image Analysis at Eindhoven University of Technology, E-mail: {r.f.p.v.pelt, m.breeuwer, b.m.terhaarromeny, a.vilanova}@tue.nl.

• Javier Oliván Bescós is with the department of iXR Innovation at Philips Healthcare, The Netherlands, E-mail: javier.olivan.bescos@philips.com.

• Marcel Breeuwer is with the department of MR Clinical Science at Philips Healthcare, The Netherlands, E-mail: marcel.breeuwer@philips.com.

• Rachel E. Clough is with the division of Imaging Sciences and the department of Vascular Surgery, NIHR Comprehensive Biomedical Research Centre of Guys and St Thomas NHS Foundation Trust and Kings College London, E-mail: rachel.clough@kcl.ac.uk.

• M. Eduard Gröller is with the department of Computer Science, within the group of Computer Graphics and Algorithms at Vienna University of Technology, E-mail: groeller@cg.tuwien.ac.at.

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.

a tedious and time-consuming task. Moreover, the segmentation process is exceptionally hard to automate for malformed morphologies. As a result, physicians often resort to local 2D segmentations, losing valuable 3D context for qualitative inspection.

We propose a framework for interactive exploration of the unsteady blood-flow field, eliminating the requirement for an a-priori segmentation of the anatomical structures. A qualitative exploration can provide sufficient information, in which case segmentation tasks are fully avoided. In case additional analysis is required, comprehensive visual exploration is likely to support identification of unforeseen regions of interest. These regions may subsequently be segmented locally.

In our framework, qualitative inspection can be performed in real-time by means of a virtual probe, inspired by operator-driven US acquisition. The virtual probe is positioned within the blood-flow field, and can be repositioned interactively. The 3D interaction is facilitated by automatically aligning the virtual probe with the average blood-flow field, along the view direction. Subsequently, the probe acts as a navigational object, providing a basis for a range of flow visualization techniques. We present hardware-accelerated particles, integral lines and integral surfaces. For the visual representation of these techniques, we employ illustrative techniques to enhance perception of the blood-flow dynamics.

In summary, the main contributions of this paper are:

- An interactive virtual probe for exploration of unsteady 4D PC-MRI blood-flow data, without a-priori segmentation of anatomical structures. Exploration in 3D is facilitated by automatically correcting the orientation of the virtual probe along the view direction, based on the average blood-flow field.
- Illustrative flow visualization techniques, including particles, integral lines and integral surfaces. All presented techniques are fully hardware-accelerated, providing interactive parametrization of the visual styles.
- An evaluation with domain experts, assessing the devised workflow for interactive exploration with the virtual probe.

2 RELATED WORK

The related work can be separated into two major categories. On the one hand, we address interactive exploration and probing of multi-valued medical data. On the other hand, we consider previous work describing flow visualization techniques, with particular interest for illustrative approaches. First, we will address the MRI blood-flow specific literature.

A select but expanding group of physicians investigate 4D PC-MRI blood-flow, studying blood-flow characteristics in relation to the progression of CVD. Current inspection of the measured blood-flow data typically includes qualitative exploration, followed by a full quantitative analysis. Generally, the inspection process is preceded by a tedious segmentation step, separating cardiovascular structures, and identifying the blood-flow regions. Amongst others, such pre-clinical studies are presented by Markl et al. [17, 18] and Uribe et al. [26].




Fig. 2. A 4D PC-MRI blood-flow data set generally consists of 20 to 25 phases in time, spanning a full cardiac cycle. Each phase in the series comprises a velocity-vector volume with a resolution of about $150 \times 150 \times 50$ voxels, sized $2.0 \times 2.0 \times 2.5$ mm. Velocity data (a) right to left (b) anterior to posterior (c) head to feet, and (d) anatomical scan.

A number of tools facilitate inspection of the 4D MRI blood-flow data. However, they are typically unable to prevent the time-consuming segmentation step. Many tools incorporate functionality for local 2D segmentation of the vasculature. This may be performed either manually, such as the contour delineation in GTFlow by Gyrotools [22], or automatically as presented in previous work [27]. Alternatively, the cardiovascular morphology may be segmented in 3D, which is typically performed using specific tools. In a scarce case, full segmentation is incorporated in the workflow, such as recently presented by Hennemuth et al. [9].

In order to avoid full segmentation, we have investigated interactive techniques, facilitating exploration of high dimensional data. In the past years, the medical visualization community has introduced a wide range of techniques, striving for simplified interaction for data exploration. In particular, Kohlmann et al. presented an extensive framework, simplifying 3D interaction, including a selective picking approach [14]. Similarly, there is the established notion of data probing, which was readily applied to flow fields by De Leeuw and Van Wijk [5]. As opposed to their glyph based probing, other approaches were proposed to locally probe data characteristics, such as the work presented by Mlejnek et al. [20]. They introduced profile flags, locally probing knee cartilage information obtained from MRI data.

The virtual probe presented in this work provides a seeding basis for a variety of flow visualization approaches. One of these approaches comprises a particle visualization. This topic is extensively described in literature. For example, hardware-accelerated approaches such as presented by Kipfer et al. [12] and Krüger et al. [16], are related to the particle system we present.

In addition, we present integral line visualizations. A large body of work describes the generation of integral lines, and for our purpose we focus on methods that allow real-time interaction. Interactive integral lines were readily presented by Mattausch et al. [19], discussing evenly-spaced streamlines. More recently, Cuntz et al. [4] and Köhn et al. [15] have presented integral line tracing techniques, employing capabilities of modern consumer graphics hardware.

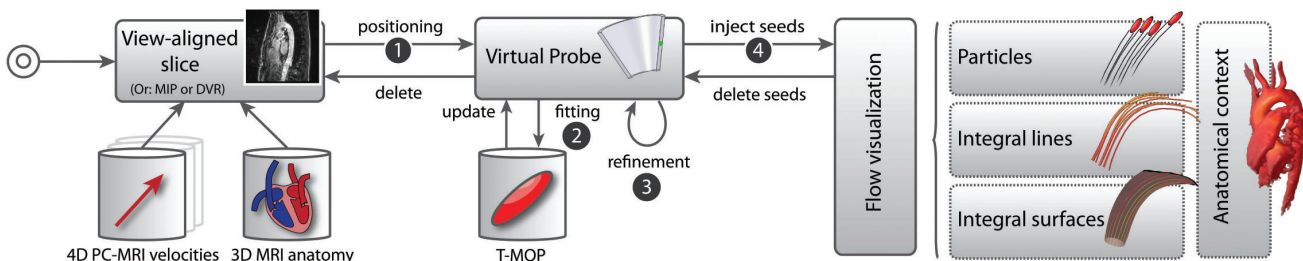


Fig. 3. A virtual probe for visual exploration of 4D PC-MRI blood-flow data. (1) Initial placement of the probe is performed on a view-aligned planar reformat (subsection 4.1). (2) Next, the probe can be aligned with the pre-computed T-MOP orientation field (subsection 4.2). (3) Subsequently, the position and orientation of the probe can be refined interactively (subsection 4.3). (4) The positioned virtual probe is represented by a half-open tapered cylinder (subsection 4.4), and provides an interactive seeding basis for various flow visualizations, including particles, integral lines and integral surfaces (section 5).

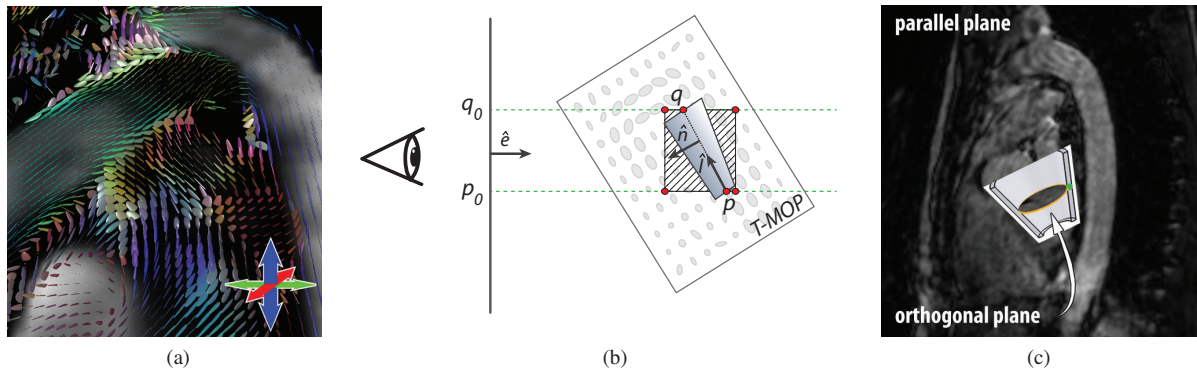


Fig. 4. (a) The T-MOP provides an average orientation field of the unsteady velocity data (see equation (2)). The orientation field is depicted by normalized ellipsoid glyphs, with RGB color-coding for XYZ orientations. The magnitude of the orientation distributions is conveyed by the gray-value intensities. (b) The virtual probe is automatically aligned to the orientation field, along view-direction \hat{e} . (c) The newly positioned virtual probe can be equipped with a parallel and orthogonal planar reformat.

Lastly, we incorporate integral surfaces. Garth et al. [7] have described a generic hardware-accelerated approach for generating pathsurfaces, while Bürger et al. [3] presented a real-time technique for the generation of streak surfaces. Recently, Born et al. [2] and Hummel et al. [10] introduced distinct illustrative visualization styles for integral surfaces. In particular, Hummel et al. [10] proposed to enhance perception of the surface shape by stripe patterns. We have adopted similar stripe patterns for our application-specific pathsurfaces.

3 BLOOD-FLOW EXPLORATION

Nowadays, US is the prevailing modality for inspection of blood-flow information. Besides its cost efficiency, US is attractive since it provides flexible and interactive imaging. A skilled operator expeditiously navigates to various regions of interest, using a manually controlled probe that locally captures blood-flow information.

This observation underpins our proposition for a probe-based workflow, enabling interactive inspection of 4D PC-MRI blood-flow measurements for pre-clinical research. As opposed to the US hardware probe, we propose a virtual probe that allows navigation within the acquired 4D PC-MRI blood-flow field.

The shape of the virtual probe was envisioned to conform to anatomical structures, and was simultaneously expected to be simple and easy to perceive. Abstraction of the cardiovascular morphology yields a cylinder as the elementary vascular shape, while a cone was determined to convey the elementary shape of a heart chamber. As a result, the probe is represented by a tapered cylinder, enabling visualization of both a cone and a cylinder.

The workflow of our framework, depicted in figure 3, entails four major steps. First, the virtual probe is positioned, based on anatomical or blood-flow information. After positioning, described in subsection 4.1, the user can directly interact with the virtual probe, including translation, rotation and scaling of the probe object.

Most common input devices, such as the ubiquitous computer mouse, allow for 2D interaction, while the virtual probe resides in a 3D environment. We intend to provide elementary 2D interactions for exploration with the probe. A method is required that automatically finds the hard-to-determine position and orientation of the probe along the view direction. We propose to align the virtual probe with the average blood-flow field, at the optimal location along the view direction. This is performed by an automated fitting approach, described in subsection 4.2.

Subsequently, the virtual probe can be refined and repositioned manually, as described in subsection 4.3. The probe is represented by a half-open tapered cylinder, as described in subsection 4.4, and serves as a seeding basis for various flow visualization approaches, which are elaborated in section 5.

4 BLOOD-FLOW VIRTUAL PROBE

4.1 Probe positioning

A new virtual probe can be added to the system using two distinct methods. On the one hand, we present a single-click ‘pen-stroke tool’. With this tool, the user creates a new line by dragging the mouse, creating a new virtual probe with the long axis identical to the user-defined line. On the other hand, we introduce a ‘two-click tool’, where the user selects the start and end points of the long axis in sequence.

Initial positioning of the probe is performed using a customary planar reformat, presenting either anatomical data or blood-flow velocity data. The planar reformat is continuously oriented towards the viewer. Different orthogonal viewpoints can be selected, mimicking conventional slice-based navigation through the data. Additionally, oblique reformats can be obtained by viewpoint interaction, using a common trackball 3D camera scheme.

Alternatively, a new probe can be positioned based on a maximum intensity projection (MIP) or a direct volume rendering (DVR). In contrast to probe positioning based on a view-aligned planar reformat, no depth information is available when positioning on a MIP or DVR. The depth information is introduced by the fitting approach, described in the next subsection.

4.2 Probe fitting

Goal: With common 2D interaction techniques, it is straight-forward to translate, rotate and scale the virtual probe in a 3D environment. However, interactive adjustment of the position and orientation along the view-direction is non-trivial. The depth and orientation of the probe are difficult to estimate visually.

We aim to reduce the freedom of interaction with the virtual probe to two dimensions. This requires an automatic detection of position and orientation of the probe, along the view direction. Therefore, we propose using the available velocity information of the unsteady blood-flow field, aligning the probe to this field. In particular, we introduce an automated fitting approach that aligns the orientation of the long axis of the virtual probe, to be tangential to the average local blood-flow orientation.

Pre-processing: Before pre-computing the average blood-flow orientations, we address the uncorrelated noise affecting the blood-flow velocity data. This undesirable spike-noise is caused by acquisition artefacts and the flow-field reconstruction process. We employ the vector median filter [1] to largely remove this noise, inducing robustness for the fitting approach. In addition, the filter moderately regularizes the velocity field without averaging, keeping tissue boundaries intact. The vector median filter VM is computed from vectors \vec{v} in a local neighborhood η surrounding the current position \mathbf{x}_i :

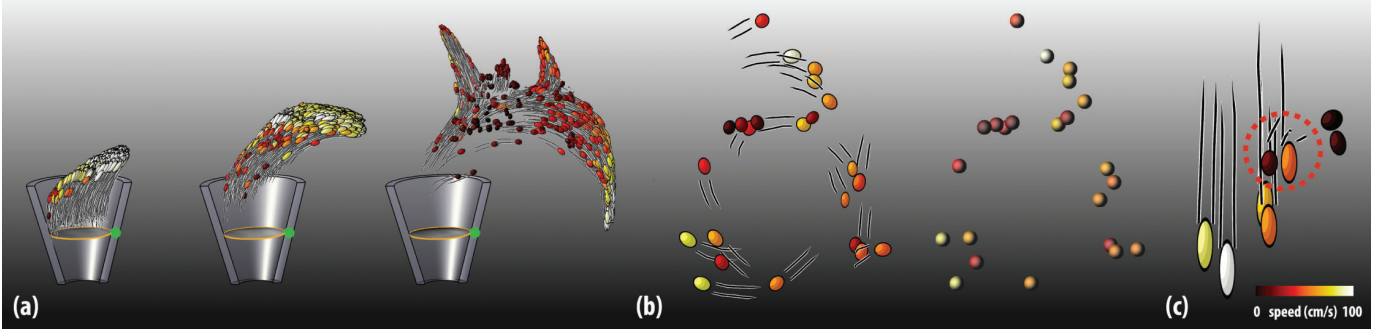


Fig. 5. Illustrative particles can be injected interactively, with the virtual probe as seeding basis. (a) Subsequently, the ‘bolus’ of particles moves through the cardiovascular system (aorta is depicted). (b) Our illustrative approach includes ellipsoidal particles with speed lines, in contrast to traditional impostor rendering of spheres. (c) Aligned speed lines indicate that the particle resides within the flow field; otherwise the particle dwells near the tissue boundary (i.e., particles within the red dashed circle).

$$VM(\mathbf{x}_i) = \bar{\mathbf{v}}(\mathbf{x}_{i_{\min}}), \text{ with } i_{\min} = \underset{i \in \eta}{\operatorname{argmin}} \left(\sum_{j \in \eta} \|\bar{\mathbf{v}}(\mathbf{x}_i) - \bar{\mathbf{v}}(\mathbf{x}_j)\| \right) \quad (1)$$

Thereafter, blood-flow orientations are averaged through time, based on the filtered velocities. The pre-computed volume, named the temporal mean-orientation-projection (T-MOP), provides a voxel-wise mean distribution of the blood-flow orientations for a full cardiac cycle. The required orientation distributions are defined by a structure tensor, as introduced by Knutsson [13]. A subset of a T-MOP volume is depicted in figure 4a, where normalized ellipsoids represent the orientation distributions.

These structure tensors are computed through the tensor product \otimes of each velocity $\bar{\mathbf{v}}$, at all voxel positions \mathbf{x} in the data set. Subsequently, the temporal mean of these tensors is computed, resulting in a description of average orientation distributions, projected over T phases of the cardiac cycle. Phases in this context are defined as time points of the acquired series, as conventionally used in the medical domain (see figure 2). The tensor-valued T-MOP volume is given as:

$$T-MOP(\mathbf{x}) = \frac{1}{T} \sum_{t=1}^T (\bar{\mathbf{v}}(\mathbf{x}, t) \otimes \bar{\mathbf{v}}(\mathbf{x}, t)) \quad (2)$$

The use of structure tensors for the T-MOP computation avoids direction cancellation, in contrast to averaging of the velocity vectors. This comes at the cost of losing directional information, which does not pose a limitation for our fitting approach. Furthermore, temporal averaging of the orientations ensures that the fitting approach is invariant to temporal fluctuations.

Objective function: Fitting the probe to the mean orientations of the unsteady velocity-field, described by the T-MOP volume, can now be defined as an optimization problem, as depicted in figure 4. First, the optimization process requires a measure to compare the orientation of the probe long axis with the T-MOP orientations, defining the objective function of the optimization problem.

To that end, coherence between orientations provides a suitable measure, defining the degree of alignment between orientation distributions in terms of direction and magnitude. In particular, we adopt the eigenvalue coherence (EC) measure, introduced by Persson et al. [21]. They define coherence as a ratio between the ordered eigenvalues λ_1 and λ_2 of the structure tensor \mathbf{T} as:

$$EC(\mathbf{T}) = \left(\frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2} \right)^2 \quad (3)$$

Using the eigenvalue coherence measure, the degree of alignment between an arbitrary vector $\hat{\mathbf{u}}$ and the T-MOP orientation at a position \mathbf{x} is then defined as the point coherence (PC):

$$PC(\mathbf{x}, \hat{\mathbf{u}}) = EC(T-MOP(\mathbf{x}) + (\hat{\mathbf{u}} \otimes \hat{\mathbf{u}})), \text{ with } \hat{\mathbf{u}} = \frac{\bar{\mathbf{u}}}{\|\bar{\mathbf{u}}\|} \quad (4)$$

Subsequently, we integrate the point coherence PC along the probe long axis to find the total line coherence (LC). The line coherence represents the alignment between the T-MOP orientation field and the probe long axis, defined by the two outermost points \mathbf{p} and \mathbf{q} . The total line coherence provides the final objective function for the optimization method, and is defined as:

$$LC(\mathbf{p}, \mathbf{q}) = \left(\frac{1}{\Delta} \sum_{i=0, \Delta}^1 PC(\mathbf{p} + i(\mathbf{q} - \mathbf{p}), \hat{\mathbf{l}}) \right), \text{ with} \quad (5)$$

$$\Delta = \frac{s}{\|\mathbf{q} - \mathbf{p}\|} \text{ and } \hat{\mathbf{l}} = \frac{\mathbf{q} - \mathbf{p}}{\|\mathbf{q} - \mathbf{p}\|}$$

The sampling distance Δ along the probe long axis is proportional to the resolution s of the T-MOP orientation field.

Optimization: Lastly, the position and orientation of optimal alignment between the probe and the T-MOP field needs to be determined. The optimization space, spanned by the lines \mathbf{p}_{δ_1} and \mathbf{q}_{δ_2} along the view direction $\hat{\mathbf{e}}$, must be searched to find the desired fit between the probe long axis and the orientation field. The probe fit optimization (PF) searches for this optimal alignment, and is defined as:

$$PF(\mathbf{p}_0, \mathbf{q}_0) = \underset{\delta_1, \delta_2}{\operatorname{argmax}} (LC(\mathbf{p}_{\delta_1}, \mathbf{q}_{\delta_2})), \text{ with } \begin{matrix} \mathbf{p}_{\delta_1} = \mathbf{p}_0 + \delta_1 \hat{\mathbf{e}} \\ \mathbf{q}_{\delta_2} = \mathbf{q}_0 + \delta_2 \hat{\mathbf{e}} \end{matrix} \quad (6)$$

As a proof of concept, we search the entire parameter space for maximum coherence between the probe long axis and the orientation field. The speed of searching the T-MOP orientation field is improved by a branch-on-need octree [28] space partitioning. This results in a local fitting time in the order of one second, which is fast enough for our needs. An additional performance improvement may be obtained by employing a faster optimization method, such as gradient descent or simulated annealing.

4.3 Probe refinement

For the envisioned qualitative exploration, we explicitly perform neither global nor local segmentation. Instead, the presented initial positioning and fitting of the virtual probe provides a plausible approximation of the desired position and orientation, for which some further refinement may be necessary.

Similar to the initial positioning, full three degrees-of-freedom for interactive refinement of the probe provides too much leeway to the target user. Instead, we restrict probe movement to a view-aligned 2D plane, parallel to the probe long axis. This restriction narrows down the interaction to a manageable two degrees-of-freedom task, without sacrificing positioning liberty. Especially in combination with the presented fitting approach, the navigation becomes much more practicable.

Within the parallel plane, the user can intuitively translate and rotate the probe. Translation is performed by dragging the object, while




Fig. 6. GPU-generated long pathlines, statically seeded in the heart chambers to capture aorta and pulmonary arteries. (a) Pathlines may be represented by illuminated lines, or (b) shaded imposter tuboids. The frames show a larger version of the pathlines, respectively without and with halos and contours. In addition, pathsurfaces are employed to convey blood-flow information. (c) We present both cross-shaped surfaces, and (d) tube-shaped surfaces. The latter may be nested, approximately capturing the wave-front profile. A stripe pattern is employed to convey rotation, (e) in contrast to conventional Phong-shaded surfaces.

rotation is enabled through a dedicated handle on the virtual probe. Additionally, the base and top radius of the virtual probe can be adjusted.

4.4 Probe visualization

Each newly positioned virtual probe is represented by a half-open tapered cylinder. Additionally, the probe can be equipped with planes parallel and orthogonal to its long axis, conveying anatomical context. The virtual probe with corresponding planes is depicted in figure 4c.

Occlusion of information within the probe is avoided by always orienting the open side of the cylinder towards the viewer. During interaction, the probe rotates about its long axis, directing the probe normal towards the viewpoint (figure 4b). Hence, the updated probe normal \hat{n} is defined as the projection of the view vector \hat{e} into the normal plane: $\hat{n} = \hat{e} - (\hat{e} \cdot \hat{l})\hat{l}$.

This approach introduces two issues. First, the view-alignment behavior is ill-defined whenever the view vector is approximately parallel to the probe long axis. In practice, this does not pose a problem, since the flow information is never occluded in this case.

Second, the proposed view alignment leads to counter-intuitive user interaction. This is due to the fact that the probe depends on the view direction, in addition to its position and orientation. Whenever the user interactively rotates the virtual probe, it is restricted to the probe parallel plane, providing basic 2D interaction. Simultaneously, the probe would be constantly updated to adhere to the view alignment, causing a conflict with the user-imposed rotation. We resolve this conflict by disabling view alignment whenever the user rotates the probe, followed by an animation that gradually corrects the resulting misalignment.

5 BLOOD-FLOW VISUALIZATION

5.1 Particle Traces

Methodology: In previous work, particle systems were readily applied to blood-flow fields. Commonly, particles are depicted as spheres [24], or otherwise represented by small integral lines [22]. Both approaches convey blood-flow speed through color, while direction information is captured by temporal cohesion. Integral lines additionally provide a short history of the particle trajectory.

Both conventional approaches employ the available visual cues, such as color and shape, to capture merely the blood-flow velocity information. Instead, we propose an illustrative particle-based approach that captures the velocity information by means of shape, keeping the color cue available for more elaborate blood-flow characteristics. To that end, we have extended the common sphere or integral line rendering. This is inspired by illustrative techniques that are typically found in cartoons. Some of our example renderings are given in figure 5.

First, spheres are deformed into ellipsoids, conveying speed and direction of the local velocity. This technique is often applied in comics, and mimics the behavior of a deforming ball in high-speed motion.

Second, we enhance the perception of direction by means of speed lines, often used to convey motion in comics [11]. We apply two lines per particle, depicting the particle trajectory for a limited amount of time. Using two lines has the advantage of capturing local reliability of the particle, indicating if the particle resides within the blood-flow stream. Whenever the two speed lines are visually aligned, the particle fully resides within the flow field. Otherwise, the particle probably dwells at a position near the tissue boundary, as emphasized by the particles within the red dashed circle in figure 5c.

The thickness of the speed lines is adjusted according to the local speed, where higher speed implies thinner lines. This choice is specific to our application, where particles in high speed blood-flow commonly group together. Thinner lines for these densely-packed particles reduce visual clutter.

The novel combination of ellipsoids and speed lines captures local blood-flow speed and direction. Consequently, color coding of speed information is redundant, saving the color cue to convey more elaborate flow characteristics. Examples are vorticity (see the leftmost rendition in figure 1) or particle residence time. A visual comparison between common sphere rendering and our illustrative approach is depicted in figure 5b.

The particle system is seeded with the virtual probe as basis. The attached orthogonal plane is typically used as a seeding plane. Alternatively, the volume spanned by the probe may be filled with seed points. For both the planar and volumetric seeding approaches, seed points may be generated according to a uniform distribution or a quasi-random Poisson-disk distribution [25]. Subsequently, interactive exploration is enabled by changing the position of the virtual probe, immediately injecting seed points at the new probe location.

Each seed point, and the associated particle movement, depends on the spatial location of the probe and the currently selected time-frame. We present a ‘seed injection’ approach, analogous to the injection of contrast agent during MRI acquisition. A set of seed points - or ‘bolus’ - can be injected interactively during animation of the particle system. Figure 5a depicts three time frames of particle traces, originating from such a bolus. Alternatively, the animation can be paused to inject seed points at a specific time during the cardiac cycle.

When particles leave the boundaries of the acquired volume, new seeds are commonly injected into the system at the original time of seeding. Alternatively, we have experimented with a continuous seeding approach, directly re-injecting seeds whenever particles vacate the volume. Each time a particle crosses the boundaries of the acquired volume, we visually fade the particle, before creating new seeds.

Implementation: The presented particle system extensively uses modern graphics processing units (GPUs), providing real-time rendering. We employ a ping-pong buffer approach, alternating particles between two vertex buffer objects (VBO). Additionally, each particle set is equipped with a second VBO, containing particle attributes such as particle age and local blood-flow velocity.

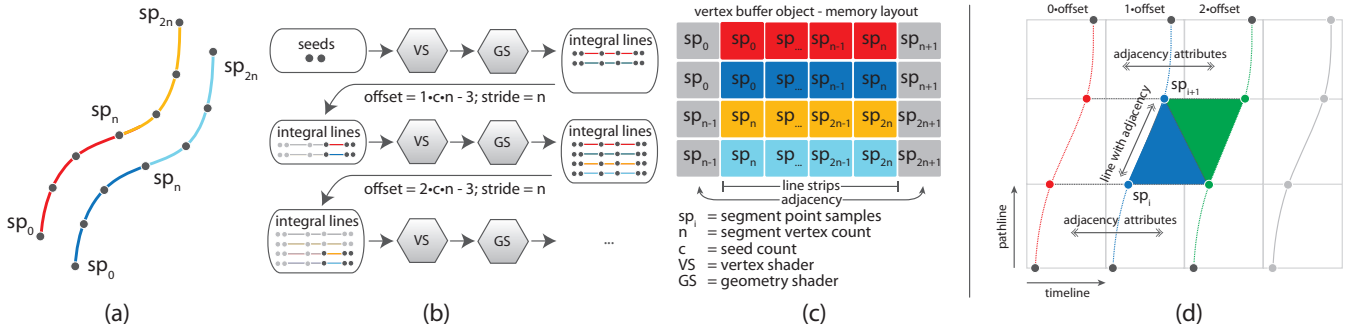


Fig. 7. (a) Overview of the GPU-driven integral-line generation. (b) Line-strip segments with adjacency information are iteratively processed by the GPU pipeline, (c) and stored to the line buffer, depending on their implicit topology. (d) Overview of the GPU-driven generation of integral-surfaces with adjacency information, exploiting parallelism in both pathline tracing and triangle construction.

The particle system is divided into three stages that are executed for each render cycle. The first stage is named the *evolution* stage, controlling aging of particles in the systems. A vertex shader with transform feedback is employed, managing particle birth and death in parallel.

Subsequently, particles are processed by the *advection* stage, displacing particles based on the unsteady velocity field. Displacement is based on a GPU-driven Runge-Kutta 4 integration scheme, approximating the spatiotemporal particle trajectory. Due to the moderate temporal resolution, extensive temporal supersampling is required.

Lastly, particles are depicted by the *rendering* stage. We employ geometry shaders for imposter rendering of the ellipsoids and associated speed lines. For each ellipsoid, a quad is generated and rotated according to the projection of the local velocity vector, as previously presented by Gumhold [8] and Prčková et al. [23]. Instead of ray-casting the ellipsoids, we use a normal-map texture, which is more efficient and delivers satisfactory results when combined with an anti-aliasing scheme.

The speed lines consist of short reverse-traced integral lines, similar to the integral lines described in subsection 5.2. Each speed line is generated in real-time by the geometry shader. The lines consist of view-aligned quad-strips extended with halos. The halos are generated during the fragment-shading stage [6].

5.2 Integral lines

Methodology: In the field of flow visualization, integral lines are commonly applied to capture the flow-field structure. In particular, pathlines are known to convey the temporal behavior of unsteady flow fields, such as the MRI measured blood-flow data.

Our framework incorporates pathlines, either represented by illuminated lines or shaded tuboids. These visual styles are depicted respectively by figures 6a and 6b. Perception of the spatial relations between pathlines is improved by means of halos [6]. Additionally, contours are applied in order to enhance the structure of the pathlines.

In order to trace the pathlines, spatiotemporal seed points are injected, similar to the particle system presented in the previous subsection. We distinguish two types of seeding. On the one hand, lines may be seeded *statically*, tracing the lines from a fixed position in space and time. On the other hand, lines may be seeded *dynamically*, tracing the lines from a fixed spatial location, and varying seed time with the current time frame of the cardiac cycle.

A statically seeded pathline is typically traced for a considerable amount of time. This approach allows inspection of the particle trajectories throughout the whole cardiac cycle. However, reliability of the lines decreases with the rapidly accumulating uncertainty of the trace.

Dynamically seeded pathlines consist of comparatively short traces [27]. Although the covered temporal range is relatively narrow, the pathlines are more reliable, and provide an approximative depiction of the pulse-wave in the cardiovascular system. The drawback of the

limited local information may be alleviated by employing multiple distinct probes.

Implementation: In previous work, various GPU-based integral-line trace approaches were proposed. For example, Köhn et al. [15] present a hardware-driven approach for fiber tracking. However, their approach requires intermediate data transfers involving GPU memory. Alternatively, Cuntz et al. [4] present integral lines, employing the geometry shader and transform-feedback. The geometry shader allows transformation of one input primitive into another one, which for line tracing amounts to a transformation from a point to a line primitive. However, Cuntz et al. observe that continuation of long line traces requires the last generated point of the previous segment, as well as the generated line segment itself. Since the geometry shader cannot return distinct primitive types, they forward the last position to the fragment-shading stage, storing the last positions to specific fragment locations.

In contrast, we present a GPU-driven approach for integral-line generation, omitting both intermediate buffer transfers and the unwieldy fragment-based storage process. We observe that vertex buffers, stored on the GPU side, may contain both geometry primitives and associated attributes, such as normal vectors and color information. Addressing these attributes requires a buffer offset, indicating the start of the specific attribute values. Instead of addressing specific attribute values, we employ the buffer-offset functionality to sequentially address segments of the line trace.

The first step of the iterative approach, depicted in figure 7a, renders the provided seed points and generates the first line segments in parallel. Each generated segment consists of the maximum number for output vertices n . Subsequently, an offset pointer is set to address the last points of the previous segment in the current line buffer, providing the seed positions for the newly traced segment. This process is iterated, increasing the segment count c , until the desired trace length is achieved. This is depicted in figure 7b.

While generating the line primitives, adjacency information between segments should be included in the generation process. We efficiently store the line segments as line strips with adjacency information. Therefore, the geometry shader is set to output point primitives, as opposed to lines. This way, the implicit topology of the lines can be manually generated by the geometry shader, resulting in a buffer memory layout as depicted in figure 7c.

Finally, the generated line-strip buffer can be rendered interactively, properly dealing with the adjacency information. For an illuminated-line representation, the tangential component is computed for each vertex, requiring the full adjacency information. Additionally, imposter rendering of tuboids requires both the tangential component, as well as the per-vertex binormal.

5.3 Integral surfaces

Methodology: In addition to particles and integral lines, we have investigated the use of integral surfaces. Pathsurfaces are generated in real-time, conveying the temporal behavior of the blood-flow. To the




Fig. 8. (a) Raycasting of the T-MIP data [27]. (b) in comparison to raycasting of anatomical data. Volume clipping is employed to avoid clutter while intermixing the presented blood-flow visualizations with a volume rendered context. (c) Either orthogonal planes (d) or the probe parallel plane may be used to cut the volume. (e) For example, illustrative particles can be visualized, together with an anatomical context that is clipped using the probe parallel plane.

best of our knowledge, the value of integral surfaces for the inspection of blood-flow measurements has not been explored before.

The pathsurfaces are illuminated according to the standard Blinn-Phong scheme, and are textured with alternating stripes. Especially for tube-shaped surfaces, the stripes emphasize the rotational motion around the centerline. The color of the stripes may convey various derived measures of the blood-flow. Additionally, the propagation front is accentuated by an outline.

The virtual probe again provides a seeding basis for the real-time surface traces. The quality of the pathsurfaces is sensitive to the placement of a sparse number of seed points. Therefore, it is essential that all seed points reside within the blood-flow field. Seed points that are close to the vessel boundaries may cause erroneous pathsurfaces, due to partial volume effects. Interactive probe refinement allows direct correction of the seed positions, whenever seeds are located outside the blood-flow regions.

Different seeding patterns lead to different integral-surface shapes. In detail, we have experimented with cross-shaped and tube-shaped pathsurfaces, as depicted in figures 6c and 6d. Additionally, multiple nested tube-shaped pathsurfaces are considered, as opposed to a single tube.

For all pathsurfaces, we propose a seeding approach similar to the one proposed for the dynamically seeded pathlines, presented in the previous subsection. Each pathsurface trace is initiated at a fixed spatial location, while the seed time is adjusted to the currently inspected time point.

Implementation: A range of approaches for integral-surface generation were presented in the literature. Garth et al. [7] presented an accurate generation of integral surfaces in unsteady flow fields, and focus on an adaptively-refined propagation of the surface front. They observe the separation between pathline and timeline generation, as depicted in figure 7d. They employ this characteristic to approximate the integral surfaces in parallel.

Similarly, we exploit this intrinsic parallelism. First, pathlines are generated in parallel, as described in the previous subsection. Subsequently, the surface is triangulated per timeline in parallel, using the implicit topology of the surface. This approach enables real-time generation of pathsurfaces, and hence facilitates interactive exploration within the blood-flow field.

Corresponding to the presented pathline traces, adjacency information should be incorporated in the generation process. In particular, triangulation and per-vertex normals require spatiotemporal information of the direct neighborhood. As depicted in figure 7d, this neighborhood information is exposed to the graphics pipeline. This is done by rendering the line primitive ($s_{p_i}, s_{p_{i+1}}$) with adjacency information, while providing neighboring lines (i.e., in red and green) as additional attributes. These line attributes are provided through appropriate buffer offsets, addressing the desired primitives within the buffer.

Various illustrative renditions of pathsurfaces were presented in recent work [2, 10]. Similar to Hummel et al. [10], we apply a stripe

pattern to enhance the surface shape. The stripes are procedurally generated in the fragment shader, allowing interactive color mapping of various derived measures of the blood-flow. The stripe pattern applied to tube-shaped surfaces is depicted in figure 6d, in contrast to conventional Phong-shaded surfaces, as depicted in figure 6e.

5.4 Anatomical context

The presented flow-visualization approaches are difficult to interpret without a notion of anatomy in the rendition. In addition to the planar reformats associated with the probe (figure 4c), we employ a customary GPU-based raycasting with geometry intermixing and interactive volume clipping.

On the one hand, the volume renderer may use blood-flow speed information, or alternatively a temporal maximum-intensity projection (T-MIP) [27], as input data. The speed information provides a cine representation of the blood-flow, while the T-MIP provides a static representation, roughly approximating the interior of the cardiovascular morphology. Both these data sets contain strong boundaries between blood-flow regions and stationary tissues. Consequently, transfer-function classification is manageable and iso-surfaces may be appropriately illuminated. An example of a volume rendered T-MIP data set is depicted in figure 8a.

On the other hand, detailed anatomical scans are often acquired. These anatomical data can similarly serve as a contextual rendering, after spatiotemporal alignment with the 4D blood-flow data. Even though the anatomical data provide much richer information, they are likewise harder to classify. Moreover, these anatomical data are typically acquired at one or two time frames during the cardiac cycle, providing a static context. An example of a volume rendered anatomical data set is depicted in figure 8b.

In order to correctly depict the spatial relations between the flow visualizations and the anatomical context, we intermix opaque geometry with semi-transparent volume raycasting. The intermixed representation, however, leads to unacceptable occlusion of the blood-flow visualizations.

Visibility of the blood-flow visualization is recovered by clipping the volume. Basic orthogonal clipping planes may be used for cutting, as depicted in figure 8c. Additionally, we employ the plane parallel to the long axis of the virtual probe as an interactive oblique clipping plane, as depicted in figure 8d. The positioned probe readily defines a region of interest, and clipping based on the probe is therefore likely to produce sensible views on the anatomy.

Intermixing of illustrative particles with a volume-rendered anatomical context is depicted in figure 8e. The anatomical data is clipped using the probe parallel plane.

6 EVALUATION

The initial conception for an exploratory virtual probe originates from discussions with physicians, actively involved with advancements in the emerging field of 4D PC-MRI blood-flow acquisition. The value

of the fully elaborated system was measured by means of an evaluation questionnaire, performed with four physicians. All respondents have in-depth general knowledge of CVD diagnosis, prognosis, and treatment. For the evaluation, we distinguish between cardiac-oriented and vascular-oriented respondents, based on the distinct ways of analyzing the blood-flow data.

Positioning: First we consider the probe-positioning approach, based on the ‘pen-stroke’ and ‘two-click’ tools, described in subsection 4.1. While the differences are seemingly small, we have come to understand that the cardiac-oriented respondents prefer orthogonal views to select end-points of the probe centerline sequentially, which is only possible with the two-click tool. In addition, all respondents preferred a reduced degree of freedom for selecting the initial oblique reformat. Alternatively, three user-defined points may define the initial probe parallel plane, based on orthogonal slice views.

Fitting: Subsequently, we have evaluated the automated fitting approach, as described in subsection 4.2. Regions of interest are commonly selected, based on anatomical scans. Our approach adds a view-dependent alignment of the probe according to the associated velocity field. This was positively valued by all the respondents, specifically for regions that are challenging to segment. Several respondents indicated that the fitting approach in practice may be used as an additional step, validating their conventional anatomy-based positioning.

All respondents expressed the need to carefully control the visual probe location. Consequently, global positioning of the virtual probe based on MIP or DVR was not particularly desired. The approximate behavior of the presented fitting approach was not regarded disadvantageous. The close alignment with the average velocity field was considered adequate, especially in combination with the interactive refinement of the probe location.

Refinement: The refinement of the virtual probe, as described in section 4.3, was evaluated particularly positive. The reduced degrees-of-freedom by means of the in-plane probe interaction was highly appreciated.

Visualization: All respondents preferred a semi-transparent glass style for the virtual probe, extended with contour lines. The continuous view alignment of the virtual probe was considered intuitive, and the associated parallel plane was greatly appreciated for anatomical context or additional blood-flow information. The orthogonal plane was valued positively for refining the probe location, even though the information is communicated from an oblique angle. Moreover, the orthogonal plane was appreciated as a seeding basis for the flow-visualization approaches.

Particles: The next part of the questionnaire addressed the various flow visualization techniques, presented in section 5. All respondents agreed that the illustrative depiction of the particles enhances the perception of the *blood-flow recirculation*. The physicians could easily relate to this representation, and value the possibility to convey other blood-flow measures by means of color. The fading cue for disappearing particles was appreciated by two of the respondents. Remarkably, three out of four respondents indicated that other physicians may mistrust the cartoon-inspired style, questioning the reliability of the conveyed information. Nevertheless, the respondents themselves clearly appreciated the novel stylistic improvements.

Besides the visual representation, also seeding of the particles was evaluated. All respondents appreciated the injection analogy, mimicking the MRI contrast-agent behavior. Three out of four respondents questioned if they would use real-time injection of seeds, and prefer to pause the animation to inject the seeds. All respondents unanimously rejected continuous seeding of particles.

Integral lines: Thereafter, the presented integral-line approaches were evaluated. In previous work, the value of integral-lines for blood-flow analysis has been readily established. In this questionnaire, we evaluate several visual styles, as well as the virtual-probe based exploration.

All respondents preferred the tuboid representation over conventional illuminated lines. Two out of four respondents, however, note that the radius of the tuboids should be only slightly larger than a typical line rendering. Halos were again valued positively for depth perception. With respect to the contour rendering, the respondents were

far from unanimous.

Altogether, the interactivity of the pathlines was considered the most valuable feature. All respondents indicated that interactive lines allow them to inspect the *blood-flow structure* in more regions within the same restricted time budget. Other tools typically require time-consuming local segmentation of vessel cross-sections and off-line pre-computation of the integral lines.

Integral surfaces: The integral surfaces for blood-flow visualization caused some controversy between the respondents. The evaluation revealed a dichotomy between vascular-oriented and cardiac-oriented physicians. The vascular-oriented physicians responded positively to the surface visualizations, indicating the benefit to capture *blood-flow rotation* in large thoracic arteries. Nesting the tubular surfaces was said to reasonably approximate the wave-front profile. In contrast, cardiac-oriented physicians questioned the value of integral surfaces over the previously presented integral lines. They argue that within the heart chamber it is hard to define regions with rotation about a longitudinal axis.

Context: The last aspect of the visualization comprises the volume rendering of the anatomical context. All respondents appreciated the presented volume rendering, in particular based on the T-MIP [27] or blood-flow speed information. However, they prefer to use an anatomical scan for contextual information, involving tedious and time-consuming transfer-function classification. Therefore, planar-reformats currently suffice as anatomical context for general purposes, while volume rendering may occasionally serve as an aesthetically-pleasing broad overview.

Finally, all respondents positively valued clipping. In particular, orthogonal planes are appreciated for their simplicity. Using probe-based clipping was valued positively as a concept, and becomes worthwhile when the current interaction with the oblique clipping plane is getting less involved. In general, the presented GPU-based volume rendering with clipping was valued over opacity modulated volume renderings, presented in competitive tools.

The results of the evaluation questionnaire are presented in table 1. In this table, ‘+’ indicates that the majority of physicians valued the aspect positively, while ‘-’ indicates a predominantly negative judgment. This is also encoded in colors, respectively in green and red, while orange (or \pm) indicates diverging judgments.

virtual probe	positioning	+	two-click tool, ortho-plane navigation
		\pm	pen-stroke tool
		-	view-aligned plane navigation
	fitting	+	concept, accuracy and speed
blood-flow visualization		+	planar reformat as basis
		-	MIP or DVR as basis
	refinement	+	translation, rotation, scaling
	visualization	+	view-alignment, associated planes
particles		+	ellipsoids, speedlines with halos
		+	contours, interactive seeding
		\pm	fading cue for leaving particles
		-	continuous seeding
integral lines		+	lines, tuboids with halos, interactivity
		+	static and dynamic seeding
		\pm	contours
	integral surfaces	+	tubes with nesting, stripe pattern
context		-	cross, only vascular oriented
		+	planar-reformat, raycasting, clipping
		\pm	probe-based clipping, T-MIP data
		-	transfer-function

Table 1. Summary of the evaluation questionnaire

7 RESULTS

The presented exploratory system was developed using the C++ programming language, supported by the OpenGL graphics library, and the visualization toolkit (VTK). In addition, the QT library was employed to construct the user interface.

The exploratory character of the system heavily depends on the attained interactivity. The GPU is employed for the majority of the visualization approaches and related algorithms, allowing real-time parametrization and interaction. The graphics hardware should at least implement the unified shading architecture.

We have carried out a quantitative performance evaluation, using a conventional system with dual-core processor, 6 GB internal memory and an NVidia GeForce 570GTX graphics card. For all renditions, a multi-sampling scheme was applied to avoid aliasing. The imposter particle visualization renders at a minimum of 40 frames per second (fps). Adding particles did not cause a noticeable performance decrease. In contrast, the performance of the pathline rendering decreases linearly with the number of lines. Interactive frame rates are obtained for up to 700 lines. The pathsurfaces are by default generated from 40 seed points, providing a smooth surface. A single tube is rendered at about 45 fps, while three nested tubes render at 30 fps. For these results, rendering of the probe and the T-MIP based context was included. The volume rendering can put a heavy load on the performance, especially with detailed anatomical data.

We have also evaluated computation times of the probe-fitting approach. The global fitting of the probe depends on the length of the ray-volume intersection. The fitting on average completes within two seconds, when octree optimization is applied. Instead, the local fitting generally finishes within one second. Using the octree, this can be reduced to merely dozens of milliseconds. In practice, increasing the octree depth by one or two levels yields satisfactory computation times and fitting results.

The fitting approach requires a T-MOP volume, which takes about ten seconds to compute. However, including median smoothing increases this computation time to twenty minutes. This time would reduce significantly with a parallel implementation.

Several 4D MRI blood-flow data sets were explored, according to the probing approach proposed in this work. Throughout the paper, several figures were obtained from a volunteer data set (figure 2, 4a, 4c, 5, 6a and 6b). The rightmost rendition in figure 1 is based on this volunteer data set as well.

Moreover, the figures were obtained from two patient studies, inspecting an aortic dissection. In particular, the first three renditions in figure 1 convey regions in the pathological blood-flow field of the first patient data set. The context visualizations in figure 8 are obtained from the same data. Figures 6c-e were obtained from a second patient study. The patient studies were approved by the local research ethics committee (study no. 08/H0809/49).

8 DISCUSSION AND CONCLUSIONS

In conclusion, we have presented a framework that facilitates interactive exploration of 4D MRI blood-flow information. In particular, we introduce a virtual probe as navigational object, providing a real-time seeding basis for various illustrative blood-flow visualizations.

All interactions with the virtual probe are reduced to two degrees-of-freedom, apart from conventional camera positioning in the 3D scene. The hard-to-determine position and orientation of the virtual probe along the view direction is detected using an automated fitting approach. The presented system incorporates various interactive flow-visualization approaches, inspired by illustrative techniques.

Both the virtual probe and the presented flow visualizations were evaluated with four physicians, who are actively involved with 4D MRI blood-flow acquisition. All respondents positively valued the virtual probe, although interaction with the initial view-aligned planar reformat is still considered somewhat cumbersome. Interaction with the virtual probe was perceived as intuitive. The presented fitting approach was greatly appreciated as a valuable addition to anatomy-based navigation.

In general, the visualization approaches were considered complementary, where particles are used to capture blood-flow recirculation, while pathlines capture temporal structures of the blood-flow. The pathsurfaces were considered valuable to detect local rotations, which proved to be of interest to vascular-oriented physicians. To the best of our knowledge, pathsurfaces were not previously applied to measured blood-flow. Therefore, we have focussed on evaluating the use of pathsurfaces for blood-flow analysis, provisionally omitting surface refinement. The quantitative evaluation shows that all components perform at interactive frame rates. This was greatly appreciated, and considered a necessary aspect for exploration.

The framework facilitates local inspection of the blood-flow information, limiting the scope to a specific spatiotemporal region. However, physicians have a priori knowledge of the anatomical regions to inspect, based on experience with 2D-cine blood-flow acquisition. Interactive local exploration around the customary regions-of-interest easily provides a substantial improvement, while global inspection remains a major visualization challenge.

The ability to explore the data without a pre-defined segmentation enables physicians to obtain a quick initial insight into the blood-flow dynamics. Based on the initial findings, the physician may judge on the acquisition quality, or may proceed with a thorough quantitative analysis. This analysis can now be performed locally, speeding up segmentation and computation of relevant blood-flow parameters.

Due to the lack of segmentation, combined with integration inaccuracies, visual artefacts will occur. The integration primitives may cross morphological boundaries, because of the limited temporal resolution and partial volume effects. With this in mind, the physicians prefer to keep the outliers visible for pre-clinical research purposes. A conventional velocity threshold to eliminate outliers is not appropriate, especially in the case of aortic dissection patients, where regions of slowly flowing blood are the primary regions-of-interest. It would be worthwhile to quantify these inaccuracies, and possibly devise improvements.

In the future, other applications of the generic virtual-probe concept may be investigated. Thusfar, the virtual probe is employed as seeding basis for a variety of flow visualizations. However, other types of probes may be defined. For example, we envision separate clipping probes, and probes that refine or cut segmentation surfaces. Furthermore, probes may be used to locally approximate quantitative information, such as vessel radii. However, an accurate quantitative analysis will require a local segmentation, in which case the probe could facilitate navigation to specific regions-of-interest.

Furthermore, different visualization techniques for the anatomical context may be investigated. While the presented volume rendering provides satisfactory results, proper transfer-function specification is still a time-consuming task.

ACKNOWLEDGMENTS

The volunteer and patient 4D MRI blood-flow data presented in this article was provided courtesy of the division of Imaging Sciences, King's College London at St Thomas' hospital. In particular, we express our sincere gratitude and appreciation to professor E. Nagel, professor T. Schaeffter, Dr. P. Beerbaum, and Dr. I. Valverde. Furthermore, we thank the vis-group at the institute of Computer Graphics and Algorithms of the Vienna University of Technology for their valuable feedback. The work presented in this publication is partially supported by the Austrian Science Fund (FWF) grant no. TRP 67-N23 (KASI).

REFERENCES

- [1] J. Astola, P. Haavisto, and Y. Neuvo. Vector median filters. *Proceedings of the IEEE*, 78(4):678–689, April 1990.
- [2] S. Born, A. Wiebel, J. Friedrich, G. Scheuermann, and D. Bartz. Illustrative stream surfaces. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1329–1338, 2010.
- [3] K. Bürger, F. Ferstl, H. Theisel, and R. Westermann. Interactive streak surface visualization on the GPU. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1259–1266, November 2009.

- [4] N. Cuntz, A. Pritzkau, and A. Kolb. Time-adaptive lines for the interactive visualization of unsteady flow data sets. *Computer Graphics Forum*, 28(8):2165–2175, 2009.
- [5] W. C. de Leeuw and J. J. van Wijk. A probe for local flow field visualization. In *Proceedings of the 4th conference on Visualization '93*, IEEE Visualization, pages 39–45. IEEE Computer Society, 1993.
- [6] M. H. Everts, H. Bekker, J. B. Roerdink, and T. Isenberg. Depth-dependent halos: Illustrative rendering of dense line data. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1299–1306, 2009.
- [7] C. Garth, H. Krishnan, X. Tricoche, T. Bobach, and K. Joy. Generation of accurate integral surfaces in time-dependent vector fields. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1404–1411, November 2008.
- [8] S. Gumhold. Splatting illuminated ellipsoids with depth correction. In *Vision, Modeling, and Visualization (VMV)*, pages 245–252, 2003.
- [9] A. Hennemuth, O. Friman, C. Schumann, J. Bock, J. Drexler, M. Huellebrand, M. Markl, and H.-O. Peitgen. Fast interactive exploration of 4D MRI flow data. *Proceeding of SPIE*, 7964:79640E–1 – 79640E–11, 2011.
- [10] M. Hummel, C. Garth, B. Hamann, H. Hagen, and K. Joy. Iris: Illustrative rendering for integral surfaces. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1319–1328, November 2010.
- [11] A. Joshi and P. Rheingans. Illustration-inspired techniques for visualizing time-varying data. In *IEEE Visualization 2005*, pages 679 – 686, Los Alamitos, CA, USA, 2005. IEEE Computer Society.
- [12] P. Kipfer, M. Segal, and R. Westermann. ÜberFlow: a GPU-based particle engine. In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS conference on Graphics hardware*, HWWS '04, pages 115–122, New York, NY, USA, 2004. ACM.
- [13] H. Knutsson. Representing local structure using tensors. In *The 6th Scandinavian Conference on Image Analysis*, pages 244–251, Oulu, Finland, June 1989.
- [14] P. Kohlmann, S. Bruckner, A. Kanitsar, and M. E. Gröller. Livesync: Deformed viewing spheres for knowledge-based navigation. *IEEE Transactions on Visualization and Computer Graphics*, 13:1544–1551, November 2007.
- [15] A. Köhn, J. Klein, F. Weiler, and H.-O. Peitgen. A GPU-based fiber tracking framework using geometry shaders. In M. I. Miga and K. H. Wong, editors, *Medical Imaging 2009: Visualization, Image-Guided Procedures, and Modeling*, volume 7261, pages 72611J–1 – 72611J–10. SPIE, 2009.
- [16] 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, November 2005.
- [17] M. Markl, M. T. Draney, M. D. Hope, J. M. Levin, F. P. Chan, M. T. Alley, N. J. Pelc, and R. J. Herfkens. Time-resolved 3-dimensional velocity mapping in the thoracic aorta: Visualization of 3-directional blood flow patterns in healthy volunteers and patients. *Journal of computer assisted tomography*, 28(4):459–468, 2005.
- [18] M. Markl, P. J. Kilner, and T. Ebbers. Comprehensive 4D velocity mapping of the heart and great vessels by cardiovascular magnetic resonance. *J Cardiovasc Magn Reson*, 13:7, 2011.
- [19] O. Mattausch, T. Theußl, H. Hauser, and M. E. Gröller. Strategies for interactive exploration of 3d flow using evenly-spaced illuminated streamlines. In K. Joy, editor, *Proceedings of Spring Conference on Computer Graphics*, pages 213–222. SCCG, April 2003.
- [20] M. Mlejnek, P. Ermes, A. Vilanova, R. van der Rijt, H. van den Bosch, F. Gerritsen, and M. E. Gröller. Profile flags: a novel metaphor for probing of T2 maps. In *IEEE Visualization 2005*, pages 599–606, Los Alamitos, CA, USA, 2005. IEEE Computer Society.
- [21] M. Persson, J. E. Solem, K. Markenroth, J. Svensson, and A. Heyden. Phase contrast MRI segmentation using velocity and intensity. *Scale Space and PDE Methods in Computer Vision*, 3459(4):119–130, March 2005.
- [22] K. Pruessmann, P. Boesiger, S. Kozerke, and U. Dydak (Gyrotools). GT-Flow. <http://www.gyrotools.com/products/gt-flow.html> (last visited).
- [23] V. Prčkovska, T. H. J. M. Peeters, M. van Almsick, B. ter Haar Romeny, and A. Vilanova. Fused DTI/HARDI visualization. *IEEE Transactions on Visualization and Computer Graphics*, 99, 2010.
- [24] S. Venkataraman (Nvidia). 4D visualization of cardiac flow. http://www.nvidia.com/content/GTC-2010/pdfs/2009_GTC2010.pdf. GPU technology conference, 2010.
- [25] J. S. Sobel, A. S. Forsberg, D. H. Laidlaw, R. C. Zeleznik, D. F. Keefe, I. Pivkin, G. E. Karniadakis, P. Richardson, and S. Swartz. Particle flurries: Synoptic 3D pulsatile flow visualization. *IEEE Comput. Graph. Appl.*, 24:76–85, March 2004.
- [26] S. Uribe, T. Tangchaoren, V. Parish, I. Wolf, R. Razavi, G. Greil, and T. Schaeffter. Volumetric cardiac quantification by using 3D dual-phase whole-heart MR imaging. *Radiology*, 248(2):606–614, June 2008.
- [27] R. van Pelt, J. O. Bescos, M. Breeuwer, R. E. Clough, M. E. Gröller, B. ter Haar Romenij, and A. Vilanova. Exploration of 4D MRI blood flow using stylistic visualization. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1339–1347, 2010.
- [28] J. Wilhelms and A. Van Gelder. Octrees for faster isosurface generation. *ACM Transactions on Graphics*, 11:201–227, July 1992.
- [29] World Heart Organization. Fact sheet cardiovascular diseases. <http://www.who.int/mediacentre/factsheets/fs317/>, 2011.