







Parallel Compositing of Volumetric Depth Images for Interactive Visualization of Distributed Volumes at High Frame Rates

Aryaman Gupta^{1,2,3} , Pietro Incardona^{1,2,3}, Anton Brock¹, Guido Reina⁴ , Steffen Frey⁵ , Stefan Gumhold¹ ,
Ulrik Günther^{6,2,3} , and Ivo F. Sbalzarini^{1,2,3} 

¹Technische Universität Dresden, Faculty of Computer Science, Dresden, Germany

²Max Planck Institute of Molecular Cell Biology and Genetics, Dresden, Germany

³Center for Systems Biology Dresden (CSBD), Dresden, Germany

⁴VISUS, University of Stuttgart, Stuttgart, Germany

⁵University of Groningen, Groningen, The Netherlands

⁶Center for Advanced Systems Understanding (CASUS), Görlitz, Germany

Abstract

We present a parallel compositing algorithm for Volumetric Depth Images (VDIs) of large three-dimensional volume data. Large distributed volume data are routinely produced in both numerical simulations and experiments, yet it remains challenging to visualize them at smooth, interactive frame rates. VDIs are view-dependent piecewise constant representations of volume data that offer a potential solution. They are more compact and less expensive to render than the original data. So far, however, there is no method for generating VDIs from distributed data. We propose an algorithm that enables this by sort-last parallel generation and compositing of VDIs with automatically chosen content-adaptive parameters. The resulting composited VDI can then be streamed for remote display, providing responsive visualization of large, distributed volume data.

CCS Concepts

• *Computing methodologies* → *Distributed algorithms; Rendering*; • *Human-centered computing* → *Visualization techniques*;

1. Introduction

Interactive volume rendering is commonly used when analyzing large scalar field data generated by scientific simulations or experimental measurement devices. Rendering at high, consistent frame rates and low latency is crucial for enabling smooth viewpoint changes and zooming, which are important for gaining depth perception and spatial understanding. Distributed compute clusters are therefore commonly used to accelerate the rendering of large data, distributing the data and parallelizing the calculations among processors. But consistently high frame rates are difficult to achieve in such a setting, due to the time-consuming raycasting procedure and the remote rendering setup, which introduces network latency.

Here, we propose the use of view-dependent piecewise-constant representations of volume data, also known as Volumetric Depth Images (VDIs) [FSE13], to decouple interactive viewpoint changes and zooming from network latency and distributed volume raycasting. VDIs are generated by dividing the volume-rendering integral along each ray into chunks that store accumulated color and opacity. The resulting representation can be much smaller than the volume data [FSE13], can be compressed and streamed efficiently [FFSE14], and recent work has shown that it can be rendered at high frame rates, providing high-fidelity approximations near the

viewpoint from which it was generated [GGI*23]. However, there currently exists no method for generating VDIs on distributed volume data.

We therefore present an algorithm for sort-last parallel generation of VDIs on distributed data. In this approach, VDIs are generated on each processing element (PE) on its volume domain in parallel—we call these “sub-VDIs”—and composited in parallel to a single VDI with load-balancing in image space. We design the present algorithm such that it can adapt to arbitrary, potentially non-convex data decompositions, as may arise, for example, in *in situ* visualization of distributed computer simulations.

We benchmark the proposed method on real-world datasets. We test the parallel compositing algorithm for accuracy and scalability, showing that it can be used to enable responsive visualization at high frame rates for large distributed volume data. We provide our implementation as part of the open-source visualization library *scenery* [GPG*19]. In summary, we contribute the following:

- We propose the use of view-dependent piecewise-constant volume representations, such as VDIs, for interactive visualization of distributed volumes at high, consistent frame rates.
- We provide an efficient parallel algorithm for scalable sort-last generation of VDIs over distributed data.

2. Related Work

Before presenting the parallel VDI generation and compositing algorithm for distributed-memory volume data, we review the state of the art and related works in relevant areas and recall the main VDI concepts.

2.1. Distributed Volume Rendering

Volume rendering is widely used for the visualization of 3D scalar fields. Soon after the volume raycasting algorithm was first presented by Levoy [Lev88], parallel volume rendering began to receive research interest [Neu93, MPHK93] with the purpose of achieving interactive visualization by distributing the data and parallelizing the rendering calculations. Later progress in parallel volume rendering enabled efficient rendering at high degrees of parallelism and for large data sizes [PYR*09, HBC12].

A commonly used strategy for parallel volume rendering is sort-last rendering [MCEF94, PYR*09, HBC12]. There, the volume data are distributed among the n PEs taking part in the rendering. Each PE performs front-to-back raycasting on its data, producing a full-resolution sub-image. The sub-images from the different PEs are then composited to a single image of the overall data.

Cavin et al. [CMF05] provided a theoretical comparison of algorithms used for compositing the sub-images. Perhaps the most straightforward is the direct-send algorithm [Neu93], where the image is divided among the n PEs such that each is responsible for compositing $1/n$ -th of the total pixels in the final image. For this, each PE receives fragments of images from all other PEs, corresponding to the part of the final image that it “owns”. Peterka et al. [PYR*09] used the direct-send approach to demonstrate scalability of parallel volume rendering on an IBM BlueGene/P system.

Other frequently used compositing algorithms include the binary-swap algorithm [MPHK94], which uses a tree structure with pairs of processes communicating at every node of the tree, and the hybrid radix- k compositing algorithm [PGR*09], which blends the direct-send and binary-swap approaches offering configurable parameters for optimization on different hardware architectures. Previous work has aimed to optimize compositing by dynamically scheduling communication of sub-images for better overlap with computation [GKH16], reducing communication cost using hybrid OpenMP/MPI parallelism and GPU-Direct RDMA [GPC*17], and by compressing image data on the GPU before compositing [LMPM21]. Interactive frame rates, however, still require a finely granular domain decomposition to reduce rendering times, and responsive visualization remains a challenge due to network latency between the user and the parallel rendering cluster.

2.2. Explorable Image Representations

For responsive remote visualization, different explorable image representations have been proposed in order to decouple user interactions from rendering. Shade et al. [SGHS98] introduced the view-dependent Layered Depth Image (LDI) representation, storing multiple pixels along each line of sight. Stone et al. [SSS16] used omnidirectional stereoscopic images, rendered on remote compute clusters and reprojected locally, to enable Virtual Reality

(VR) visualization of molecular dynamics simulations. These approaches, however, are limited to surface and geometry data. For reprojecting volume data, Zellmann et al. [Zel21] used a single depth layer transmitted by the rendering server together with the color buffer, proposing various heuristics to create the depth buffer. While using a single depth value per pixel ensures small message sizes, it is not conducive to producing high-quality reprojections, as holes occur where rays do not intersect the depth layer. This has been addressed by view-dependent piecewise-constant volume representations [BJNN97, FSE13, LRBR16], which produce a continuous representation of the volume by storing multiple layers with composited color and opacity in-between. One such representation is the VDI, as described in more detail in Sec. 2.4.

Tikhonova et al. [TCM10a, TCM10b] and, more recently, Rapp et al. [RPD22], proposed compact view-dependent representations that enable interactive transfer function changes. The VDI differs from these representations in that it is generated for a given transfer function and stores transfer-function classified color and opacity within each segment. This prioritizes fast rendering from new viewpoints. Rapp et al. [RPD22] also support viewpoint changes, but require slower bilinear interpolation of Lagrange multipliers.

Recent years have seen increased research interest in novel-view synthesis—using one or more images of a scene to generate images from new viewpoints—using deep-learning techniques. Mildenhall et al. [MST*21] proposed the NeRF (Neural Radiation Fields) representation, training a neural network to encode a continuous volume in its weights, enabling rendering by raycasting over samples collected from the network. This has been extended to implicit neural representations achieving high compression ratios on large volume data [LJLB21, HW22]. While techniques have been proposed to accelerate neural rendering [MESK22, WBDM22], dense regions in large volumes still require many samples, limiting frame rates.

Exploratory visualization of numerical simulation results was done post-hoc using the Cinema database [AJO*14], which stores images generated *in situ* for a range of visualization parameters, including different viewpoints. All parameter ranges, including viewpoints, however, must be specified in advance, and the database becomes large if many viewpoints are required. Our approach instead generates new VDIs at regular time intervals, which can be streamed to enable approximate rendering with full six-degrees-of-freedom viewpoint changes.

2.3. Generation of View-Dependent Piecewise-Constant Volume Representations

We review techniques that have been proposed for generating view-dependent piecewise-constant volume representations, including VDIs [FSE13].

All view-dependent representations of volume data are generated by raycasting through the volume and decomposing the volume-rendering integral into segments, each of which containing transfer-function classified composited color and opacity. The distinguishing feature of these representations, in comparison to other techniques that compress volume data, is that they produce an exact image when rendered from the original viewpoint of generation, owing to the associativity of the `OVER` operator [PD84] used in alpha-

compositing. Rendering from deviating viewpoints involves accumulation over the segments, which is much cheaper than evaluating the full integral [KM05]. Close approximations to volume rendering are achieved around the viewpoint of generation [FSE13].

Previous works on generating view-dependent volume representations differ in the strategies used to determine the locations and extents of the segments generated along the rays. Brady et al. [BJNN97] used constant-size segments along each ray, creating segments that contain composited color and opacity over potentially heterogeneous samples, hampering the quality of rendering from a new viewpoint. Lochmann et al. [LRBR16] created segments of constant opacity by partitioning the total transmittance equally among the segments. This, however, does not account for potentially varying color values within the segments. Frey et al. [FSE13] proposed the VDI, which uses homogeneity as a criterion for segment termination, accumulating samples into a segment unless it differs from the current segment by more than a user-defined sensitivity parameter. Recent work [GGI*23] has proposed an iterative search strategy to automatically determine the sensitivity parameter. However, while VDIs can provide high-fidelity approximate renderings, there exists so far no method for generating them in parallel on distributed volume data.

2.4. The Volumetric Depth Image (VDI)

Frey et al. [FSE13] proposed the VDI as a view-dependent representation of volume data. They call the segments generated along each ray *supersegments*. Each supersegment \mathbb{S} is represented by its front and back faces, $f(\mathbb{S})$ and $b(\mathbb{S})$, and its color and opacity, $C(\mathbb{S})$ and $\alpha(\mathbb{S})$, respectively.

Each ray (x, y) cast into the volume creates a supersegment list \mathbb{L}_{xy} of up to $\mathbb{N}_{\mathbb{S}}$ (user parameter) supersegments \mathbb{S}_j^{xy} , where j is the index of the supersegment within the list (Fig. 1). The total number of lists created, $\mathbb{N}_{\mathbb{L}}$, corresponds to the viewport resolution the VDI is generated for, i.e. $\mathbb{N}_{\mathbb{L}} = wh$, where w is the width of the viewport and h the height in pixels.

The decomposition of the volume rendering integral into supersegments is governed by a termination criterion τ , which depends on a sensitivity parameter γ . Samples along each ray are merged into a supersegment until

$$\tau : \gamma < \|C(\mathbb{S})\alpha(\mathbb{S}) - C'\alpha'\|_2, \quad (1)$$

where C' and α' are the color and the length-adjusted opacity of the next sample. In words, the next sample is composited into the current \mathbb{S} unless it differs from the premultiplied color of \mathbb{S} by more than γ , in which case a new \mathbb{S} is started. This criterion generates homogeneous \mathbb{S} , which is important for generating high-quality approximated renderings from new viewpoints.

The value of γ that generates the most accurate VDI depends on the dataset and on the transfer function. A greedily optimal per-ray value is found automatically by iterative bisection search [GGI*23]. At each iteration, a pass is performed through the volume, and the total number of times τ is true is used to modify the value of γ for the next iteration given a $\mathbb{N}_{\mathbb{S}}$ budget. If τ was true more often than $\mathbb{N}_{\mathbb{S}}$, the value of γ is increased for the next iteration, and vice versa. This process iterates until γ has converged to a

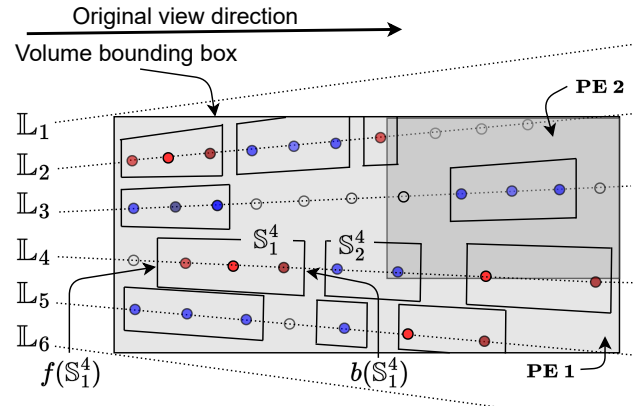


Figure 1: A *Volumetric Depth Image (VDI)* [FSE13] is generated by casting rays through the volume and grouping samples (depicted by colored circles along the ray) of similar color and opacity, generating a list \mathbb{L}_i of up to $\mathbb{N}_{\mathbb{S}}=3$ (in the example of the figure) supersegments \mathbb{S}_j^i per ray. Each \mathbb{S}_j^i stores its front and back face, $f(\mathbb{S}_j^i)$ and $b(\mathbb{S}_j^i)$, along with color and opacity accumulated in-between (Fig. 2a). Hollow circles represent samples in empty regions. Volume data may be divided among multiple Processing Elements (PE) in a computer cluster (background gray levels).

value that maximizes the number of supersegments generated while staying below a total of $\mathbb{N}_{\mathbb{S}}$. This iterative per-ray γ optimization enables VDIs to be generated without manual intervention, accurately adapting to the data and the transfer function.

In the following, we extend the VDI concept to distributed parallel generation and compositing. More specifically, we propose a compact memory layout for VDIs in order to reduce the communication overhead during parallel compositing, we show how γ search can be approximated without global communication, and we handle arbitrary non-convex data-domain decompositions. All design decisions are rationalized by performance measurements.

3. Measurement Setup

We compare and evaluate design decisions and algorithmic approaches in a uniform benchmark setup. All measurements are taken on the *taurus* high-performance computer of TU Dresden. Each node contains 8 NVIDIA A100-SXM4 GPUs with 40 GB of DRAM each, 2 AMD EPYC 7352 CPUs with 24 cores each, 1 TB RAM, and runs RedHat Enterprise Linux version 7.9. C++ code was compiled using GCC 10.3.0, Java code was run using OpenJDK 11.0.2, and OpenMPI version 4.1.1 was used. Processes are always distributed in a block manner, i.e., all 8 GPUs on a node are occupied before starting to use another.

The datasets used for the measurements are described in Table 1. All of them are commonly used for the evaluation of visualization tools and algorithms and are referred to by their name or abbreviation. Volume data is decomposed across PEs in blocks of as close to equal size and extents along the spatial dimensions as permitted by the data dimensions. Rendering of VDIs, performed to verify the

Dataset (Abbreviation)	Dimensions	Size
Kingsnake	1024×1024×795 8bit uint	795 MiB
Rayleigh-Taylor Instability [CCM04]	1024×1024×1024 16bit uint	2 GiB
Richtmyer-Meshkov [CDD*02]	2048×2048×1920 8bit uint	7.5 GiB
Rotating Stratified Turbulence [RPMM15] (RS)	4096×4096×4096 16bit uint	128 GiB
Forced Isotropic Turbulence [YDS12] (FI)	4096×4096×4096 16bit uint	128 GiB

Table 1: Datasets used for the measurements.

quality of VDIs generated, ran on a Nvidia RTX 4090 on a remote workstation under Ubuntu 20.04.

4. Compact VDI Representation

VDI generation techniques [FSE13] have so far used a regular 3D representation for VDIs. Any lists that pass through empty regions, and therefore do not generate supersegments, or that generate less than N_S supersegments, store zeros in the remaining locations as illustrated in Fig. 2a. This regular 3D structure of the VDI improves the performance of raycasting-based rendering as it leads to better memory access patterns. However, for the parallel compositing of VDIs, which requires communication of sub-VDIs between PEs, this generates unnecessary communication overhead.

For the purpose of parallel compositing, we therefore store VDIs in memory using a compact representation. Only those supersegments that are actually generated are stored, as illustrated in Fig. 2c. The difference between the memory required for a compact representation and a regular representation grows when more PEs are used in a sort-last parallel generation approach. This is because data becoming divided more finely among PEs leads to sub-VDIs that are increasingly sparse.

The procedure for generating compact VDIs starts similarly to that for regular VDIs. For each ray, a value of γ is determined using iterative search (Sec. 2.4). No supersegments are actually generated at this stage, but value of γ and the number of supersegments that it would generate are both stored in a buffer. Next, a prefix sum is calculated on this buffer, which records for each list \mathbb{L}_i the total number of supersegments generated by all lists before \mathbb{L}_i (Fig. 2b). Entry i in the prefix buffer, corresponding to list \mathbb{L}_i , therefore contains the index in the linearized array at which the supersegments of \mathbb{L}_i start to be stored. All lists can therefore generate their supersegments in parallel, using the value of γ determined in the first step. In our implementation, both γ search and supersegment generation are parallelized across lists on the GPU, while the prefix sum is calculated on the CPU.

Table 2 compares the time to generate a compact VDI representation with the time to generate a regular representation, which avoids the prefix-sum computation and the GPU kernel synchronization at the end of the per-ray γ search. We find that generating

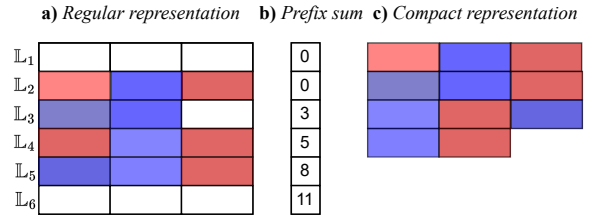


Figure 2: Representing the VDI generated in Fig. 1 in memory. a) All 18 \mathbb{S} are stored in memory. b) Prefix sum evaluated on (a), which is used to generate the compact representation of the VDI shown in c), storing only the 11 non-empty \mathbb{S} .

Dataset	Regular	Compact	Memory ratio
Kingsnake	0.42 s	0.29 s	1:0.12
Rayleigh-Taylor	0.54 s	0.44 s	1:0.45
Richtmyer-Meshkov	0.91 s	0.84 s	1:0.34

Table 2: Time in seconds to generate a single $N_L=1920 \times 1080$, $N_S=25$, VDI stored either using the regular or the compact representation in memory, and the ratio of the memory (regular:compact) required for the representations.

the compact representation is, in fact, faster than generating the regular representation, with the additional time to compute the prefix sum more than amortized by the fact that empty supersegments do not need to be written to memory.

The compact VDI representation is a form of run-length encoding for the VDI. Though a typical run-length encoding approach would group all neighboring supersegments with the same color and depth values, not only empty ones, it is unlikely in practice for neighboring supersegments to have the same value, unless they are empty. This form of lossless compression is analogous to active-pixel encoding [Mor11, LMPM21] commonly performed in parallel compositing of images, where empty regions in images are compressed using run-lengths.

Although it is possible to render VDIs in their compact representation, this is significantly slower than rendering the regular representation as it offers less optimal memory access patterns (Fig. 6). The VDI we finally stream for rendering after parallel compositing is therefore stored in the regular representation, but we use the compact representation during parallel generation and compositing.

5. Generating a VDI on distributed data

We propose a method to generate VDIs that represent data distributed across PEs, e.g., compute nodes in a cluster, GPUs within a node, etc. The final VDI represents the entire volume data in the viewport across all PEs and can be streamed for remote display.

The proposed approach relates to techniques commonly used to generate images from distributed data. We adopt a sort-last parallel rendering strategy [MCEF94] in order to achieve scalability with data size and to conform to arbitrary domain decompositions, as, e.g., produced by an *in situ* numerical simulation. Distributed VDI

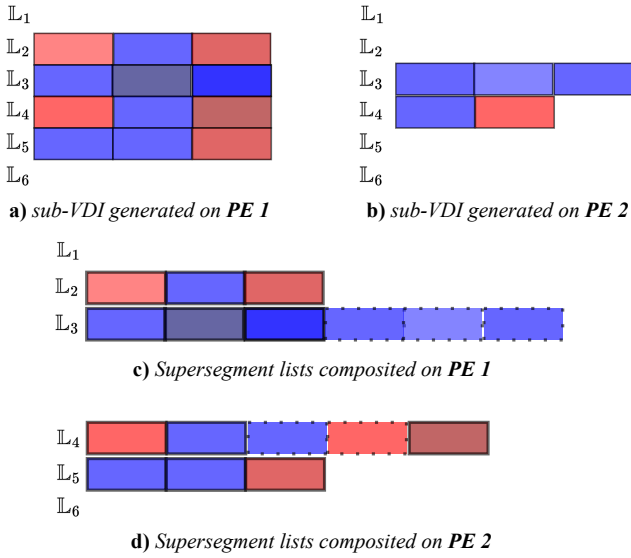


Figure 3: The sub-VDIs generated during Phase 1 and the lists composited during Phase 2 for the VDI in Fig. 1. Panels c) and d) show sub- \mathbb{S} sorted by their position along the ray. Flat outlines represent sub- \mathbb{S} received from PE 1, while dotted outlines represent sub- \mathbb{S} received from PE 2. These are then composited, producing a maximum of $N_{\mathbb{S}} = 3$ \mathbb{S} per list, as shown in Fig. 1.

generation therefore begins from a given domain decomposition where each PE only stores a part of the overall volume (Fig. 1). Each PE then generates a sub-VDI at full viewport resolution for its local data and stores it in the compact representation. The sub-VDIs are composited into a single VDI representing the entire volume data using a direct-send approach [Neu93] with compositing load balanced in image space. The compositing stage receives supersegments from all PEs and combines them to produce a total of up to $N_{\mathbb{S}}$ supersegments per list, minimizing loss of detail.

5.1. Phase 1: Distributed generation of sub-VDIs

Distributed sub-VDI generation starts from a domain decomposition of the volume data. As is typical for sort-last approaches, no transfer of data between PEs is required. A VDI corresponding to the full viewport resolution is generated concurrently on each PE. All PEs share the camera viewpoint from which rays are cast to generate supersegments. The sub-VDIs are stored in the compact representation. This avoids a linearly growing communication overhead with increasing numbers of PEs, since the sub-VDIs become increasingly sparse for finer data decompositions.

Any given ray in the view frustum can, in general, pass through the domains of multiple PEs. The search for a value of γ that would generate a total of up to $N_{\mathbb{S}}$ supersegments along the ray would thus require communication between PEs at each iteration. This would prevent the algorithm from scaling, as γ search typically requires several iterations to converge. We avoid this communication and synchronization by proposing an approximate algorithm that allows each PE to act independently during sub-VDI generation. For this,

PEs do not make assumptions about the numbers of supersegments generated on other PEs. Instead, each PE independently generates up to $N_{\mathbb{S}}$ supersegments (we refer to them as sub- \mathbb{S}) within its own domain. This ensures that the volume rendering integral is never under-resolved for the given budget of $N_{\mathbb{S}}$. In Fig. 1, for example, ray 2 passes through the domains of both PE 1 and PE 2, but encounters only empty regions on PE 2. All $N_{\mathbb{S}} = 3$ supersegments are thus generated on PE 1 (Fig. 3). Ray 3, however, generates a full $N_{\mathbb{S}} = 3$ on each PE, over-resolving the ray with six sub- \mathbb{S} , which are then reduced during compositing.

The proposed strategy of generating up to $N_{\mathbb{S}}$ sub- \mathbb{S} on each PE a ray passes through generates more sub- \mathbb{S} than required. This increases the amount of data to be communicated during compositing. But it eliminates the cost of communicating and synchronizing all PEs during γ search and preserves the best-possible quality for the final VDI.

5.2. Phase 2: Parallel compositing of sub-VDIs

At the end of Phase 1, each ray has produced up to $N_{\mathbb{S}}$ sub- \mathbb{S} on each PE where it intersects the data. These need to be merged to produce a total of up to $N_{\mathbb{S}}$ supersegments for each ray. The first step is therefore to bring all sub- \mathbb{S} of a ray from all PEs onto a single PE where they can be composited.

We design an algorithm based on the direct-send approach for compositing sub-images in distributed volume rendering [Neu93]. The number of supersegment lists \mathbb{L} in the final composited VDI is divided equally among PEs, with each PE responsible for producing composited \mathbb{L}_{xy} for the pixels in its part of the image space. In Fig. 3, e.g., both PE 1 and PE 2 composite three \mathbb{L} each. This balances the load in image space.

Each PE sends to all PEs, including itself, the sub- \mathbb{S} those PEs are responsible for compositing. Since the amounts of data to be sent to different PEs differ, we use `MPI_AlltoAllv`, which accounts for variable message lengths. The prefix-sum buffers generated for the sub-VDIs are also transmitted, in chunks corresponding to the image space decomposition, as they are required for reading sub- \mathbb{S} from the compact representation. Each PE then has all the data required for compositing the final supersegments \mathbb{S} of the lists in its part of the image space.

The first step in merging the sub- \mathbb{S} is to determine the order in which they lie along the ray. The sub- \mathbb{S} in any list cannot be assumed to be contiguous. There may be gaps between consecutive sub- \mathbb{S} when the ray passes through the domain of another PE. In Fig. 1, e.g., Ray 4 passes through the domain of PE 2 before returning to PE 1, generating sub- \mathbb{S} on PE 2 (see Fig. 3) that are to be placed in-between the sub- \mathbb{S} of PE 1. Within each PE, however, the sub- \mathbb{S} are sorted, since they are generated by front-to-back ray-casting. Therefore, to iteratively determine the next sub- \mathbb{S} in a list, the depths of the frontmost sub- \mathbb{S} from all PEs are compared, and the sub- \mathbb{S} with the lowest starting depth is selected as the next along the ray. The merged set of sub- \mathbb{S} for Ray 4 is produced on PE 2, which is responsible for compositing \mathbb{L}_4 (Fig. 3d).

The process of compositing the merged sub- \mathbb{S} can be formulated as another supersegment generation task, performed by raycasting

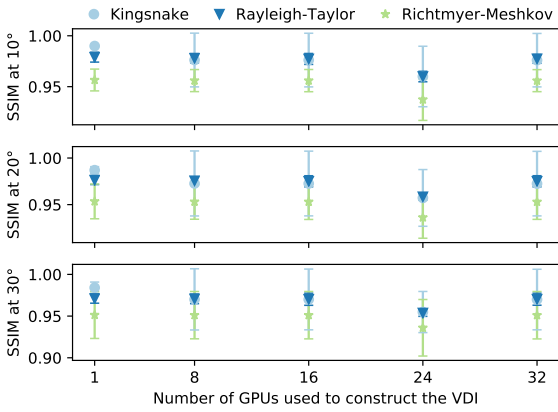


Figure 4: SSIM (mean \pm min-max range across four viewpoints of VDI generation) with respect to ground truth DVR for the rendering of VDIs generated using varying numbers of Nvidia A100 GPUs. VDIs generated on 8, 16, 24 and, 32 GPUs are composited using the presented compositing algorithm (Sec. 5.2), while VDIs generated on 1 GPU do not undergo compositing. All VDIs are of resolution $N_L=1920 \times 1080$ with $N_S=25$ for three datasets (symbols, top legend) and three different V_N (panels).

through the sub- \mathbb{S} , which are, after all, piecewise constant samples of the original volume. We therefore treat each sub- \mathbb{S} as a sample along the ray. These samples are raycast through and combined into supersegments \mathbb{S} . Since the sub- \mathbb{S} have different lengths, the process of raycasting through them is analogous to volume raycasting with irregular step size. The opacity obtained from a sub- \mathbb{S}_j^i is the opacity stored in sub- \mathbb{S}_j^i , corrected by its length [EHK*04] as:

$$\tilde{\alpha} = 1 - (1 - \alpha)^l \quad (2)$$

where $\tilde{\alpha}$ is the adjusted opacity, α is the opacity stored in sub- \mathbb{S}_j^i , and l is the length of sub- \mathbb{S}_j^i . Empty spaces between sub- \mathbb{S} are treated as transparent samples. At each sample, the sub- \mathbb{S} can either be merged with the previous supersegment, or it can begin a new one. This is again determined using the criterion τ (Eq. 1) and requires finding another γ that leads to the generation of N_S supersegments in total. This is again done using per-ray iterative γ search (Sec. 2.4), which requires multiple passes through the sub- \mathbb{S} . The sub- \mathbb{S} in \mathbb{L}_4 in Fig. 3d are, e.g., combined to produce the \mathbb{S} depicted in Fig. 1. Detailed pseudocode is included in the Supplement.

While parallel compositing approaches like binary-swap [MPHK93] and radix-k [PGR*09] typically outperform the direct-send for image compositing, we choose the latter here as it requires only a single stage of compositing sub- \mathbb{S} into \mathbb{S} . The proposed sort-last approach of compositing sub- \mathbb{S} produces different \mathbb{S} depending on the number of PEs the data is divided over. To evaluate the accuracy, we generate VDIs on multiple GPUs and compare the quality of the rendering with a VDI generated on a single GPU, where compositing is not required. VDIs are generated from four viewpoints (V_O) representing 90° rotations of the dataset in camera space. They are then rendered at different viewpoint deviations (V_N) about the viewpoint of generation, and quality is compared to

ground truth direct volume rendering (DVR). Figure 4 presents the results with image similarity measured using the SSIM [WBSS04] metric, depicting the mean and the range (max-min) across the four V_O . While there is a marginal decline in SSIM values, we observe that VDI rendering quality remains similarly high for VDIs generated on multiple GPUs as for a VDI generated on a single GPU. Transfer functions are depicted in Fig. 5 and images are included in the Supplement for visual comparison.

Figure 6 compares the rendering frame rates of VDIs stored in regular and compact representation with the DVR baseline at different V_N about V_O . The regular representation renders faster due to its simpler memory-access patterns. The task of converting from compact to regular representation cannot be handled by the display client as it needs to remain responsive for user interaction. Therefore, while the sub-VDIs communicated for compositing used the compact representation to reduce communication overhead, the supersegment lists produced by compositing use the regular representation. The composited lists from all PEs are combined onto the root PE by `MPI_Gather`, from where they are streamed for (remote) display, potentially after lossless compression.

5.3. Handling Non-Convex Data Decompositions

A key feature of the above compositing method is that it can handle non-convex domain decompositions and therefore work with any application-given data distribution.

A non-convex domain decomposition is one where a ray can intersect the boundary of the domain of a PE in more than two points. Such decompositions occur, e.g., in numerical simulations in complex-shaped simulation domains, where the domain decomposition balances the computations in each sub-domain and the communication volume between PEs [ILZ*19], not necessarily producing an equal division of data among PEs. Such situations are challenging for distributed volume rendering, due to the non-commutativity of the `over` operator [PD84]:

$$a \text{ over } b \neq b \text{ over } a. \quad (3)$$

This implies that in non-convex domain decompositions, volume rendering cannot composite color across disjoint segments of a ray without requiring communication or synchronization between the PEs, or redistribution of the volume data.

Our method avoids this problem by generating sub- \mathbb{S} that store world-space front and back depths along the ray. A sub- \mathbb{S} necessarily terminates when the ray leaves the domain of a PE. Since sub- \mathbb{S} are ordered by their depth during compositing, `over` operations are done in the correct order. The sub- \mathbb{S} along a ray can therefore be generated in parallel without synchronization or communication between the PEs.

The present approach includes non-convex distributed volume rendering as a limit case: when generating only a single sub- \mathbb{S} per sub-domain intersection, the compositing algorithm can, in addition to placing the supersegments in correct order, also perform `over`-operator compositing along the supersegment lists. This effectively performs volume rendering, creating a flat image on a non-convex domain decomposition without requiring synchronization or communication between PEs.

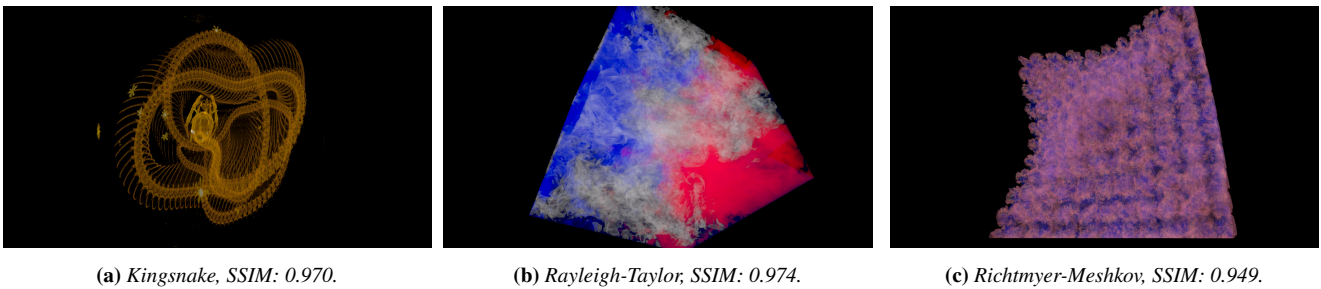


Figure 5: Visual illustration of VDI rendering quality. VDIs generated on 32 GPUs are rendered at 20° from the viewpoint of generation. SSIM values computed w.r.t. ground-truth direct volume rendering (DVR) at the same viewpoint (see Supplement for images).

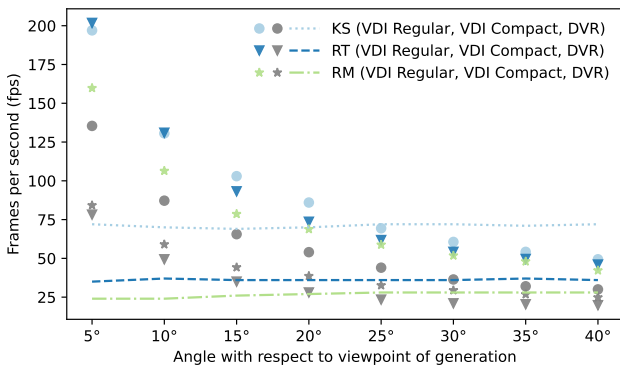


Figure 6: Performance comparison (in fps) between DVR (lines) and rendering $N_L=1920 \times 1080$, $N_S=25$, VDIs stored in either regular (colors) or compact (grayscale) representation, at various angles about the viewpoint of generation. Datasets (Kingsnake (KS), Rayleigh-Taylor (RT), Richtmyer-Meshkov (RM)) are represented by symbols/line styles, see inset legend.

6. Implementation

We implemented the algorithms described in the previous sections in the open-source rendering framework *scenery* [GPG*19]. Both sub-VDI generation and VDI compositing are implemented as compute shaders in the Vulkan API. For work distribution, a local work-group size of 16×16 is used, i.e., the screen space is divided into 2D blocks of that size. During raycasting, each ray within a block corresponds to a thread on the GPU, and to a single pixel on screen.

The final VDI generated is compressed using LZ4, which we found to provide faster and better compression than Snappy and ZSTD, before streaming. For a $N_L=1920 \times 1080$, $N_S=25$, VDI this produces ≈ 325 MiB of data, while the corresponding uncompressed VDI in regular 3D representation would be ≈ 1.2 GiB.

The source code is available under the open-source BSD license at: <https://github.com/scenerygraphics/scenery-insitu/>.

7. Benchmarks and Evaluation

We evaluate parallel VDI generation on the real-world datasets from Table 1, measuring rendering performance and quality.

We profile the performance of the individual steps of parallel compositing (Sec. 5.2) in order to quantify the benefits of using the proposed compact VDI representation. Figure 7 plots, for both compact and the regular representation, the timings of the three stages of parallel compositing: the distribution of sub- \mathbb{S} among PEs (p_d), the actual compositing into \mathbb{S} (p_c), and the final gather at the root PE (p_g). An `MPI_Barrier` is placed before all MPI calls for profiling. Measurements are reported as mean \pm standard-deviation over 144 iterations with the camera rotating 5° about the dataset every second iteration for a full 360° orbit. For p_d and p_c , timings are averaged across PEs at every iteration, since the PEs can independently proceed with the next step; p_g is recorded at the root PE. VDIs of the FI dataset (Table 1) are used in full HD (1920×1080) viewport resolution with $N_S=25$.

Using the compact VDI representation during compositing is faster in all tested cases and provides better scalability with increasing numbers of GPUs. On 32 GPUs, the entire parallel compositing process, i.e. ($p_d + p_c + p_g$), for the compact representation requires 40% of the time of the regular representation. As expected, the average number of sub- \mathbb{S} generated per PE decreases when using more PEs, as the sub-VDIs become increasingly sparse. This amplifies the advantage of the compact representation, reducing p_d by about 90% on 32 GPUs.

The reduction in the number of sub- \mathbb{S} per PE also reduces the compositing time (p_c) to $0.6 \times$ on 32 GPUs, relative to 8 GPUs. VDI sparsity, however, depends on V_O , resulting in the higher standard-deviation for the compact representation. For the regular representation, the number of sub- \mathbb{S} per PE remains constant, leading to constant compositing times. The communication volume for the final gather (p_g) is independent of the number of PEs and equal to the size of the final VDI. The observed variability is thus attributed to background load fluctuations on the benchmark machine. Eventually, p_d limits scalability, as it increases ($1.4 \times$ on 32 GPUs relative to 8 GPUs) even for the compact representation.

Figure 8 shows the overall rate of VDI generation for the 128 GiB RS [RPM15] and FI [YDS12] datasets, starting at 4 GPUs which is the minimum required for this data size (Table 1).

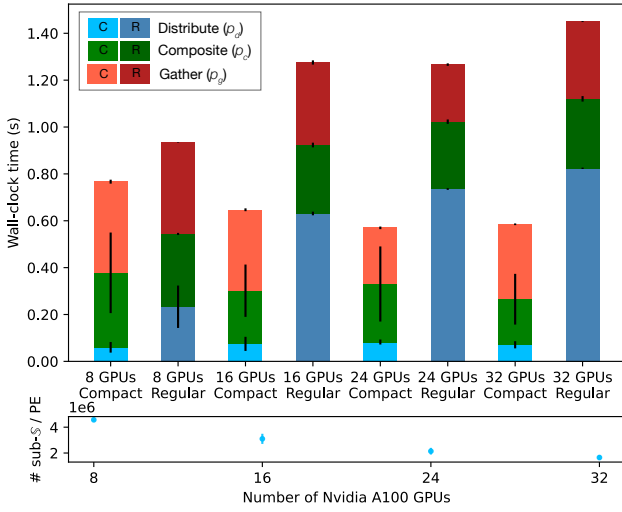


Figure 7: Runtime performance of the three stages of parallel compositing for VDIs in both compact (C, light colors) and regular (R, dark colors) representation. Mean \pm standard deviation (error bars) are reported over a 360° camera orbit (top panel). The individual parts are: distribution of sub- \mathbb{S} (ρ_d , blue), compositing into \mathbb{S} (ρ_c , green), final gathering of \mathbb{S} (ρ_g , red). The bottom panel shows the average number of sub- \mathbb{S} generated by a PE.

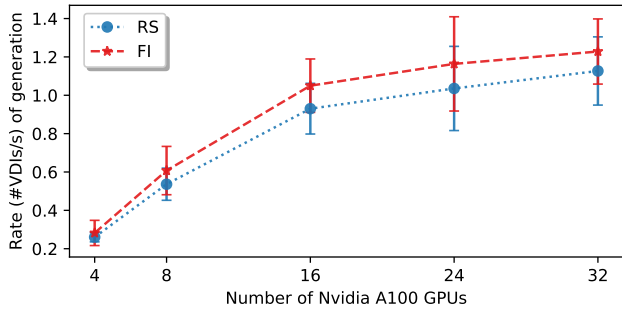


Figure 8: Number of VDIs generated per second (mean \pm standard deviation over a 360° camera orbit) for different numbers of Nvidia A100 GPUs, $N_L=1920 \times 1080$, $N_S=25$, for the RS and FI datasets of 128 GiB each (see Table 1).

Again, the camera performs a 360° orbit of the data at 5° increments with two iterations at each viewpoint.

At small numbers of GPUs, sub-VDI generation dominates the overall VDI generation time. As the number of GPUs increases, sub-VDI times decrease (strong scaling). From 4 to 8 GPUs, VDI generation rates therefore increase by $2.2\times$ and $2.1\times$ for the FI and RS datasets, respectively. As sub-VDI times decrease further, the communication overhead of parallel compositing begins to dominate, and speed-ups reduce to $1.75\times$ from 8 to 16 GPUs and $1.2\times$ from 16 to 32 GPUs. This is typical for strong scaling and is also observed in distributed volume rendering [LMPM21]. Our implementation overlaps sub-VDI generation with the final gather of the

previous VDI, effectively hiding some of the sub-VDI generation time. This increases VDI generation rates, but decreases the measured parallel efficiency. Overall, about one VDI is generated per second on 32 GPUs. The results are similar for other values of N_S . VDI generation rates increase by $\approx 15\%$ for $N_S=20$ and decrease by $\approx 9\%$ for $N_S=30$, relative to the $N_S=25$ plotted here (see Supplement).

Once the final VDI is generated, it can be streamed for interactive remote rendering. We therefore compare VDI rendering on the display client with distributed DVR on the cluster. We implement distributed DVR using IceT [Mor11], a commonly used library for sort-last parallel compositing, as also used by ParaView [AGL05] and VisIt [CBW*12]. To ensure optimal compositing performance, we recompile IceT with support for the recent CUDA image compression extension [LMPM21]. The sort-last parallel compositing strategy is chosen using the ICET_SINGLE_IMAGE_STRATEGY_AUTOMATIC option, which automatically selects between the radix-k [PGR*09], binary swap [MPHK93], and binary tree [Mor11] methods, based on the number of processes. This is the default setting in IceT.

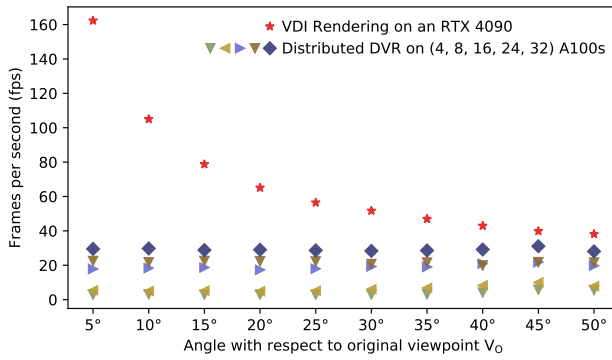
VDIs are rendered at different viewpoint deviations (V_N) about the viewpoint of generation (V_O) and compared to distributed DVR from the same viewpoint. Results are reported at four different V_O at 90° rotation increments around the data. Transfer functions with multi-hue colormaps are used to present a challenging test case for the VDI. Volume raycasting, both for DVR and VDI generation, used an emission-absorption illumination model.

Figure 9 reports rendering frame rates averaged over the four V_O . For distributed DVR, frame rates are limited by volume raycasting with parallel compositing adding an overhead of ≈ 12 ms per frame on 32 GPUs. Close to V_O , VDI frame rates are significantly higher than distributed DVR, achieving a speed-up of $5.5\times$ at $V_N=5^\circ$ for both datasets. For larger deviations, VDI frame rates reduce due to the anisotropic view-dependent shape of the VDI. But they remain at least $1.3\times$ better than for distributed DVR even at $V_N=50^\circ$.

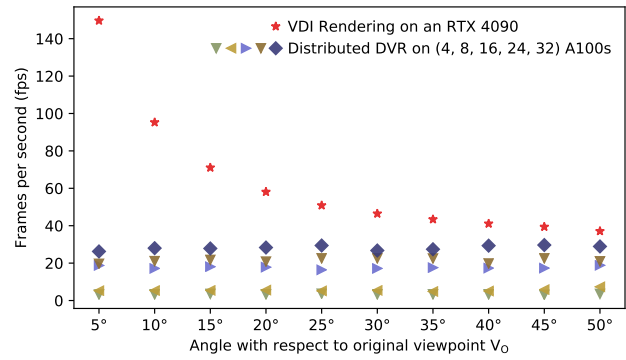
Finally, we compare the quality of the images generated by VDI rendering with those from DVR. VDIs generated on 32 GPUs at the same four V_O are rendered on a remote workstation and compared with distributed DVR on 32 GPUs. Figure 10 reports the SSIM (Structural Similarity Index Measure) [WBSS04] between the VDI rendering result and DVR, with 1.0 indicating identical images. VDI rendering provides high-quality approximations to DVR for both the RS and FI datasets. As expected, SSIM reduces with increasing viewpoint deviation, but overall remains high even at 50° . The results are similar when measuring rendering quality in terms of PSNR (Peak Signal-to-Noise Ratio, see Supplement). Visual comparisons are given in Fig. 11 with full-resolution images available in the Supplement, along with screencast videos of interactive rendering on the FI and RS datasets.

8. Conclusions

We have presented a distributed generation and parallel compositing approach for Volumetric Depth Images (VDIs), enabling responsive interactive visualization of large distributed volumes at



(a) Forced Isotropic Turbulence (FI) dataset, 128 GiB.



(b) Rotating Stratified Turbulence (RS) dataset, 128 GiB.

Figure 9: VDI rendering frame rates on a display client with one Nvidia RTX 4090 (stars), averaged over four viewpoints at 90° from each other, compared with distributed DVR on the cluster using IceT for sort-last parallel compositing on varying numbers of Nvidia A100s (other symbols, see inset legend); resolution 1920×1080 with $N_S=25$.

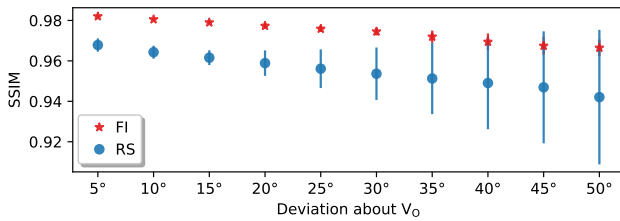


Figure 10: SSIM image similarity between VDI rendering and DVR for different V_N . Mean and min-max range (error bars) are reported over four different V_O for two datasets (RS: circles, FI: stars).

consistently high frame rates. We improved the scalability of parallel compositing using a compact memory layout for VDIs that stores only non-empty supersegments, analogous to active-pixel encoding in parallel image compositing [Mor11, LMPM21]. We found that this significantly reduced communication overhead during compositing (Fig. 7) and that VDIs in the compact representation were slightly faster to generate than those represented at regular 3D resolution (Table 2). After compositing, however, the VDIs were represented at regular 3D resolution in memory, as we found this to allow for significantly faster rendering (Fig. 6).

We followed a sort-last approach where sub-supersegments received for compositing were treated as samples of varying length and α -composited along the ray to form the final supersegments. This produced accurate VDI approximations (Fig. 4) even on large datasets (Figs. 10, 11) and allowed us to generate sub-VDIs in parallel without any communication between processing elements. This is particularly beneficial in combination with the automated γ parameter optimization, where synchronization would otherwise occur at each iteration and for each ray. Overall speed-ups for VDI generation (Fig. 8) reduced from linear to sub-linear with increasing GPU counts as soon as compositing times became dominant, as expected for strong scaling.

The proposed distributed generation approach enabled us to test the VDI on larger data than what was possible in previous works. We found that VDI rendering on a display client was about $5.5 \times$ faster than distributed DVR near the viewpoint of generation and remained faster also at large viewpoint deviations (Fig. 9). The rendered images were almost identical, particularly close to the viewpoint of generation (Figs. 10, 11).

Overall, VDIs were generated in less than one second when using 32 GPUs for a 128 GiB datasets (Fig. 8). These generation times are comparable to or lower than the iteration times of typical distributed numerical simulations. We therefore envision distributed VDI generation finding applications in interactive *in situ* visualization of numerical simulations. This also includes simulations on unstructured grids, as VDIs can be generated for any volume discretization that can be raycast. While this will influence sub-VDI generation times, as it would for any rendering, it does not change the downstream parallel compositing presented here. We further see potential applications with the Cinema database [AJO*14] for exploratory post-hoc visualization, where distributed VDIs could be used to reduce the size of the database by reducing the need for images generated from different viewpoints.

The present parallel approach for compositing sub- \mathbb{S} into \mathbb{S} maintains scalability with increasing PEs , reducing compositing times (Fig. 7) despite an increase in the overall number of sub- \mathbb{S} due to our strategy of generating up to N_S sub- \mathbb{S} per list on each PE . On the other hand, a limitation is that it is susceptible to load imbalance due to variation in the number of sub- \mathbb{S} across lists. This is because we equi-distribute lists among PEs . Future optimizations could explore alternate strategies for distributing lists that balance the sub- \mathbb{S} distribution among PEs . Whether the gain in load balance amortizes the additional global communication required to do so, however, remains to be seen. In addition, the proposed approach inherits limitations of the VDI representation. Since VDIs store transfer-function classified color and opacity, they do not support interactive transfer-function modification. Also, gradients of

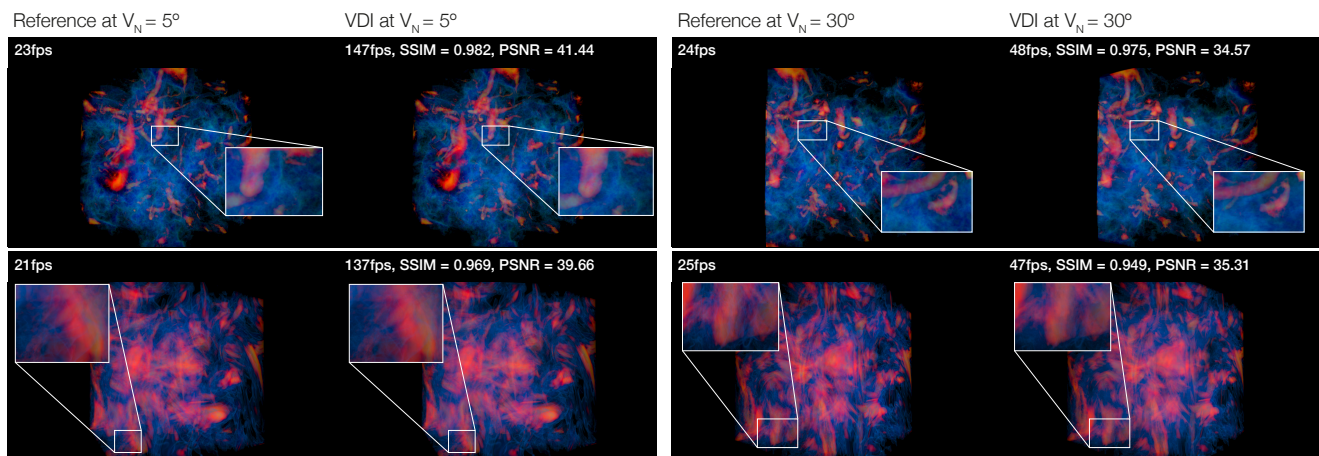


Figure 11: Visual comparison of VDI rendering quality with distributed DVR (“Reference”) for the Forced Isotropic Turbulence (top) and Rotating Stratified Turbulence (bottom) datasets with a multi-hue transfer function. Image quality metrics are computed w.r.t. the DVR image.

the volume data cannot be calculated from the VDI, precluding the use of directional illumination effects, such as specular lighting.

Notwithstanding these limitations, we believe that the methods and algorithms presented here are key in enabling the use of view-dependent volume representations, such as the VDI, for interactive visualization of large distributed volume data.

Acknowledgments

This work was supported by the Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI) Dresden/Leipzig, by the German Federal Ministry of Education and Research (BMBF) in the joint project “6G-life” (ID 16KISK001K), and by the Center for Advanced Systems Understanding (CASUS) financed by the BMBF and the Saxon Ministry for Science, Culture and Tourism (SMWK) with tax funds on the basis of the budget approved by the Saxon State Parliament. We thank the Center for Information Services and High Performance Computing (ZIH) of TU Dresden for providing their facilities for the benchmarks. We thank the University of Texas High-Resolution X-ray CT Facility (UTCT) for the Kingsnake dataset.

References

- [AGL05] AHRENS J., GEVECI B., LAW C.: ParaView: An end-user tool for large data visualization. *Visualization Handbook* (2005). 8
- [AJO*14] AHRENS J., JOURDAIN S., O’LEARY P., PATCHETT J., ROGERS D. H., PETERSEN M.: An image-based approach to extreme scale in situ visualization and analysis. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (Piscataway, NJ, USA, 2014), SC ’14, IEEE Press, pp. 424–434. doi:10.1109/SC.2014.40. 2, 9
- [BJNN97] BRADY M., JUNG K., NGUYEN H., NGUYEN T.: Two-phase perspective ray casting for interactive volume navigation. In *Proceedings. Visualization ’97 (Cat. No. 97CB36155)* (1997), pp. 183–189. doi:10.1109/VISUAL.1997.663878. 2, 3
- [CBW*12] CHILDS H., BRUGGER E., WHITLOCK B., MEREDITH J., AHERN S., PUGMIRE D., BIAGAS K., MILLER M., HARRISON C., WEBER G. H., KRISHNAN H., FOGAL T., SANDERSON A., GARTH C., BETHEL E. W., CAMP D., RÜBEL O., DURANT M., FAVRE J. M., NAVRÁTIL P.: VisIt: An end-user tool for visualizing and analyzing very large data. In *High Performance Visualization—Enabling Extreme-Scale Scientific Insight*. 10 2012, pp. 357–372. doi:10.1201/b12985. 8
- [CCM04] COOK A. W., CABOT W., MILLER P. L.: The mixing transition in Rayleigh-Taylor instability. *Journal of Fluid Mechanics* 511 (2004), 333–362. doi:10.1017/S0022112004009681. 4
- [CDD*02] COHEN R. H., DANNEVIK W. P., DIMITS A. M., ELIASON D. E., MIRIN A. A., ZHOU Y., PORTER D. H., WOODWARD P. R.: Three-dimensional simulation of a Richtmyer–Meshkov instability with a two-scale initial perturbation. *Physics of Fluids* 14, 10 (2002), 3692–3709. doi:10.1063/1.1504452. 4
- [CMF05] CAVIN X., MION C., FILBOIS A.: COTS cluster-based sort-last rendering: performance evaluation and pipelined implementation. In *VIS 05. IEEE Visualization, 2005.* (2005), pp. 111–118. doi:10.1109/VISUAL.2005.1532785. 2
- [EHK*04] ENGEL K., HADWIGER M., KNISS J. M., LEFOHN A. E., SALAMA C. R., WEISKOPF D.: Real-time volume graphics. In *ACM SIGGRAPH 2004 Course Notes* (New York, NY, USA, 2004), SIGGRAPH ’04, Association for Computing Machinery, p. 29–es. doi:10.1145/1103900.1103929. 6
- [FFSE14] FERNANDES O., FREY S., SADLO F., ERTL T.: Space-time Volumetric Depth Images for in-situ visualization. In *2014 IEEE 4th Symposium on Large Data Analysis and Visualization (LDAV)* (2014), pp. 59–65. doi:10.1109/LDAV.2014.7013205. 1
- [FSE13] FREY S., SADLO F., ERTL T.: Explorable Volumetric Depth Images from raycasting. In *2013 XXVI Conference on Graphics, Patterns and Images* (2013), pp. 123–130. doi:10.1109/SIBGRAPI.2013.26. 1, 2, 3, 4
- [GGI*23] GUPTA A., GÜNTHER U., INCARDONA P., REINA G., FREY S., GUMHOLD S., SBALZARINI I. F.: Efficient raycasting of Volumetric Depth Images for remote visualization of large volumes at high frame rates. In *2023 IEEE 16th Pacific Visualization Symposium (PacificVis)* (2023), pp. 61–70. doi:10.1109/PacificVis56936.2023.00014. 1, 3
- [GKH16] GROSSET A. V. P., KNOLL A., HANSEN C.: Dynamically scheduled region-based image compositing. In *Proceedings of the 16th Eurographics Symposium on Parallel Graphics and Visualization* (Goslar, DEU, 2016), EGPGV ’16, Eurographics Association, p. 79–88. 2

- [GPC*17] GROSSET A. V. P., PRASAD M., CHRISTENSEN C., KNOLL A., HANSEN C.: TOD-Tree: Task-overlapped direct send tree image compositing for hybrid mpi parallelism and gpus. *IEEE Transactions on Visualization and Computer Graphics* 23, 6 (2017), 1677–1690. doi:10.1109/TVCG.2016.2542069. 2
- [GPG*19] GÜNTHER U., PIETZSCH T., GUPTA A., HARRINGTON K. I., TOMANCAK P., GUMHOLD S., SBALZARINI I. F.: scenery: Flexible virtual reality visualization on the Java VM. In *2019 IEEE Visualization Conference (VIS)* (2019), pp. 1–5. doi:10.1109/VISUAL.2019.8933605. 1, 7
- [HBC12] HOWISON M., BETHEL E. W., CHILDS H.: Hybrid parallelism for volume rendering on large-, multi-, and many-core systems. *IEEE Transactions on Visualization and Computer Graphics* 18, 1 (2012), 17–29. doi:10.1109/TVCG.2011.24. 2
- [HW22] HAN J., WANG C.: Coordnet: Data generation and visualization generation for time-varying volumes via a coordinate-based neural network. *IEEE Transactions on Visualization and Computer Graphics* (2022), 1–12. doi:10.1109/TVCG.2022.3197203. 2
- [ILZ*19] INCARDONA P., LEO A., ZALUZHNYI Y., RAMASWAMY R., SBALZARINI I. F.: OpenFPM: A scalable open framework for particle and particle-mesh codes on parallel computers. *Computer Physics Communications* 241 (2019), 155–177. doi:https://doi.org/10.1016/j.cpc.2019.03.007. 6
- [KM05] KAUFMAN A., MUELLER K.: Overview of volume rendering. In *Visualization Handbook*, Hansen C. D., Johnson C. R., (Eds.). Butterworth-Heinemann, Burlington, 2005, pp. 127–174. doi:https://doi.org/10.1016/B978-012387582-2/50009-5. 3
- [Lev88] LEVOY M.: Display of surfaces from volume data. *IEEE Computer Graphics and Applications* 8, 3 (1988), 29–37. doi:10.1109/38.511. 2
- [LJLB21] LU Y., JIANG K., LEVINE J. A., BERGER M.: Compressive neural representations of volumetric scalar fields. *Computer Graphics Forum* 40, 3 (2021), 135–146. doi:https://doi.org/10.1111/cgf.14295. 2
- [LMPM21] LIPINKSI R., MORELAND K., PAPKA M. E., MARRINAN T.: GPU-based image compression for efficient compositing in distributed rendering applications. In *2021 IEEE 11th Symposium on Large Data Analysis and Visualization (LDAV)* (2021), pp. 43–52. doi:10.1109/LDAV53230.2021.00012. 2, 4, 8, 9
- [LRBR16] LOCHMANN G., REINERT B., BUCHACHER A., RITSCHEL T.: Real-time novel-view synthesis for volume rendering using a piecewise-analytic representation. In *Vision, Modeling & Visualization* (2016), Hullin M., Stamminger M., Weinkauff T., (Eds.), The Eurographics Association. doi:10.2312/vmv.20161346. 2, 3
- [MCEF94] MOLNAR S., COX M., ELLSWORTH D., FUCHS H.: A sorting classification of parallel rendering. *IEEE Computer Graphics and Applications* 14, 4 (1994), 23–32. doi:10.1109/38.291528. 2, 4
- [MESK22] MÜLLER T., EVANS A., SCHIED C., KELLER A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.* 41, 4 (July 2022), 102:1–102:15. doi:10.1145/3528223.3530127. 2
- [Mor11] MORELAND K. D.: IceT users’ guide and reference. doi:10.2172/1005031. 4, 8, 9
- [MPHK93] MA K.-L., PAINTER J., HANSEN C., KROGH M.: A data distributed, parallel algorithm for ray-traced volume rendering. In *Proceedings of 1993 IEEE Parallel Rendering Symposium* (1993), pp. 15–22. doi:10.1109/PRS.1993.586080. 2, 6, 8
- [MPHK94] MA K.-L., PAINTER J., HANSEN C., KROGH M.: Parallel volume rendering using binary-swap compositing. *IEEE Computer Graphics and Applications* 14, 4 (1994), 59–68. doi:10.1109/38.291532. 2
- [MST*21] MILDENHALL B., SRINIVASAN P. P., TANCİK M., BARRON J. T., RAMAMOORTHY R., NG R.: NeRF: Representing scenes as neural radiance fields for view synthesis. *Commun. ACM* 65, 1 (dec 2021), 99–106. doi:10.1145/3503250. 2
- [Neu93] NEUMANN U.: Parallel volume-rendering algorithm performance on mesh-connected multicomputers. In *Proceedings of the 1993 Symposium on Parallel Rendering* (New York, NY, USA, 1993), PRS ’93, Association for Computing Machinery, p. 97–104. doi:10.1145/166181.166196. 2, 5
- [PD84] PORTER T., DUFF T.: Compositing digital images. *SIGGRAPH Comput. Graph.* 18, 3 (1984), 253–259. doi:10.1145/964965.808606. 2, 6
- [PGR*09] PETERKA T., GOODELL D., ROSS R., SHEN H.-W., THAKUR R.: A configurable algorithm for parallel image-compositing applications. In *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis* (New York, NY, USA, 2009), SC ’09, Association for Computing Machinery. doi:10.1145/1654059.1654064. 2, 6, 8
- [PYR*09] PETERKA T., YU H., ROSS R., MA K.-L., LATHAM R.: End-to-end study of parallel volume rendering on the IBM Blue Gene/P. In *2009 International Conference on Parallel Processing* (2009), pp. 566–573. doi:10.1109/ICPP.2009.27. 2
- [RPD22] RAPP T., PETERS C., DACHSBACHER C.: Image-based visualization of large volumetric data using moments. *IEEE Transactions on Visualization and Computer Graphics* 28, 6 (2022), 2314–2325. doi:10.1109/TVCG.2022.3165346. 2
- [RPMM15] ROSENBERG D., POUQUET A., MARINO R., MININNI P. D.: Evidence for Bolgiano-Obukhov scaling in rotating stratified turbulence using high-resolution direct numerical simulations. *Physics of Fluids* 27, 5 (2015), 055105. doi:10.1063/1.4921076. 4, 7
- [SGHS98] SHADE J., GORTLER S., HE L.-W., SZELISKI R.: Layered Depth Images. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1998), SIGGRAPH ’98, Association for Computing Machinery, p. 231–242. doi:10.1145/280814.280882. 2
- [SSS16] STONE J. E., SHERMAN W. R., SCHULTEN K.: Immersive molecular visualization with omnidirectional stereoscopic ray tracing and remote rendering. In *2016 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)* (2016), pp. 1048–1057. doi:10.1109/IPDPSW.2016.121. 2
- [TCM10a] TIKHONOVA A., CORREA C. D., MA K.-L.: Explorable images for visualizing volume data. In *2010 IEEE Pacific Visualization Symposium (PacificVis)* (2010), pp. 177–184. doi:10.1109/PACIFICVIS.2010.5429595. 2
- [TCM10b] TIKHONOVA A., CORREA C. D., MA K.-L.: An exploratory technique for coherent visualization of time-varying volume data. *Computer Graphics Forum* 29, 3 (2010), 783–792. doi:https://doi.org/10.1111/j.1467-8659.2009.01690.x. 2
- [WBDM22] WU Q., BAUER D., DOYLE M. J., MA K.-L.: Instant neural representation for interactive volume rendering, 2022. URL: https://arxiv.org/abs/2207.11620, doi:10.48550/ARXIV.2207.11620. 2
- [WBSS04] WANG Z., BOVIK A., SHEIKH H., SIMONCELLI E.: Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* 13, 4 (2004), 600–612. doi:10.1109/TIP.2003.819861. 6, 8
- [YDS12] YEUNG P. K., DONZIS D. A., SREENIVASAN K. R.: Dissipation, enstrophy and pressure statistics in turbulence simulations at high Reynolds numbers. *Journal of Fluid Mechanics* 700 (2012), 5–15. doi:10.1017/jfm.2012.5. 4, 7
- [Zel21] ZELLMANN S.: Remote volume rendering with a decoupled, ray-traced display phase. In *Smart Tools and Apps for Graphics - Eurographics Italian Chapter Conference* (2021), Frosini P., Giorgi D., Melzi S., Rodolà E., (Eds.), The Eurographics Association. doi:10.2312/stag.20211479. 2